Prof. Barry Appleton, FCIArb
Co-Director and Distinguished Senior Fellow, Center for International Law, New York Law School
Faculty Director, ABA TechCred Program |
barry.appleton@nyls.edu

17 July 2026

European Commission
Directorate-General for Communications Networks, Content and Technology
European AI Office

Dear Sir or Madam,

Re: Targeted consultation on the draft guidelines for the classification of high-risk AI systems. The application of Annex III, point 8(a) to international arbitration.

I write in support of the draft guidelines, and to offer refinements that would make them clearer and more durable for international arbitration.

I am an international arbitration practitioner and a legal scholar. At New York Law School, I serve as Co-Director and Distinguished Senior Fellow of the Center for International Law, where I am also a Distinguished Adjunct Professor of Law. I serve as Faculty Director of the ABA International Law Section's TechCred credential in digital and AI competency for dispute resolution. These are my own views. I file them in an individual academic capacity.

Arbitration belongs within the classification. Annex III, point 8(a) reaches artificial intelligence used in a similar way in alternative dispute resolution, and Recital 61 confirms the intent where the outcome produces legal effects for the parties. An arbitral award produces legal effects. Paragraph 426 of the draft Guidelines reads alternative dispute resolution to include arbitration, and it is right to do so. The Commission should keep that conclusion and state it plainly. Certainty on this threshold serves everyone who must comply.

The harder question is not whether arbitration is covered. It is which uses are. The answer is functional.

A system that drafts the reasoning, applies the law to the facts, or identifies authorities and proposes how to apply them operates in the high-risk zone. A system that extracts metadata, searches, proofreads, or translates does not. This distinction is already in the Act. Paragraph 70 of the draft Guidelines makes clear that human involvement does not change what a system is intended to do, and Article 6(3) filters out systems that perform a narrow procedural task or that only improve a completed human activity. Classification should follow the intended purpose and the function, not the marketing of the tool. A general-purpose assistant does not become high-risk merely because an arbitrator could open it. It becomes high-risk when it is intended for, or is put to and relied upon in, the adjudicative function. Read this way, the perimeter is proportionate. It catches the risk and leaves the ordinary work of arbitration alone.

The classification should attach to the right actor. In arbitration the person who uses the system in the adjudicative function is the arbitrator, and the arbitrator is therefore the deployer. Paragraphs 428 and 429 rightly note that the Act does not define which bodies are alternative dispute resolution bodies. That definition should turn on function. Where an institution provides administrative and organizational services, such as case management, scheduling, and the logistics of appointment, it is not exercising the adjudicative function, and it should not carry deployer obligations for those services. On this point I support the request of the arbitral institutions. It is correct, and the Guidelines should grant it.

One case needs care. An institution can do more than administer. It can build and operate a system that itself performs the adjudicative function. The AAA-ICDR AI Arbitrator, launched in late 2025 outside of the EU, drafts awards for a trained human to review before signature. Where an institution builds such a system, it is the provider, and the human who adopts the output is the deployer. Each carries its own duties along the value chain. The lesson is not that institutions are suspect. It is that the line must be drawn by function, so that responsible innovation is recognised rather than exempted by accident. A human-in-the-loop design of that kind is a model to reward. A classification keyed to the label of the actor, rather than to the function performed, would exempt the very case that point 8(a) exists to reach.

Point 8(a) reaches the AI used in the adjudicative function. It does not reach the AI used by counsel or by the parties in presenting their case. That use is governed by the rules of professional responsibility and by the courts, and not by this classification. Where such use bears on the fairness of the proceeding, the concern is met through the tribunal's own duty of oversight, and not by treating counsel's tools as high-risk.

Human involvement should not defeat the classification where it is nominal. Paragraph 70 already treats oversight as a condition of compliance, not as an escape from classification. Arbitration should hold to that. Genuine judgment is required, not a signature on a machine's output. The point is not abstract. In April 2026 the Superior Court of Quebec annulled an award in ARIHQ c. Sante Quebec (2026 QCCS 1360) because artificial intelligence had generated authorities that did not exist, and the award was built upon them. The vice was the delegation of the decision, not the use of a tool. A clear functional line protects the enforceability of awards. It does not threaten it.

The draft Guidelines leave a gap on territorial scope, and it should be filled. The Act reaches deployers established in the Union, and deployers outside the Union where the output of the system is used in the Union. For arbitration the Guidelines do not say which connecting factor governs. I propose a plain answer. The relevant place for deployer status is the arbitrator's establishment or habitual residence, and not the seat, because the seat is a legal construct and not the place where a person uses a tool. Where the output of the system materially shapes reasoning that is then relied upon or enforced in the Union, that output is used in the Union within Article 2(1)(c). The Commission should resolve the reach of that test, rather than leave it to be litigated.

The line between assistance and substitution is the workable test for the grey zone. Uses that supply or shape the reasoning fall within the classification. Uses that assist without entering the reasoning do not. This tracks Article 6(3), and it can be operated as tiered disclosure, matched to what a tool actually did. Article 86 reinforces the point. A person affected by a decision taken on the basis of a high-risk system is entitled to a clear explanation of the role the system played. In arbitration that supports disclosure to the parties at the level of the proceeding.

Investment dispute bodies deserve the same functional treatment. They exercise an adjudicative function and produce awards with legal effects, often for States as well as for private parties. Paragraph 429 raises their characterization. The answer is the principle already stated. Classify by function, and the label of the forum does the least work.

The credible course for arbitration is not to seek a place outside a classification that plainly fits. It is to help the Commission draw the line well. A function-based classification protects the enforceability of awards, gives providers and deployers a clear rule, and rewards the institutions that innovate with care. I would welcome the opportunity to assist the Commission further, in writing or in a meeting, and I remain at your disposal.

Yours faithfully,

Prof. Barry Appleton, FCIArb

Co-Director and Distinguished Senior Fellow, Center for International Law, New York Law School

Faculty Director, ABA TechCred Program

Scholar and Fellow, Balsillie School of International Affairs

---