RESPONSE TO THE CONSULTATION ON THE PROPOSED IMPLEMENTING REGULATION ON DETAILED ARRANGEMENTS FOR THE CONDUCT OF CERTAIN PROCEEDINGS BY THE COMMISSION PURSUANT TO REGULATION (EU) 2024/1689

April 2026

The Association for Computing Machinery (ACM) is the world's longest-established professional society of individuals involved in all aspects of Computing. ACM's Europe Technology Policy Committee (Europe TPC) is charged with, and committed to, providing policymakers and the public with sound technical information to support sound public policymaking. ACM and Europe TPC are non-profit, non-political, and non-lobbying organizations.

Europe TPC supports the European Commission's intent to refine the regulatory framework to ensure legal certainty and safety. To help bridge the gap between legislative intent and technical feasibility, Europe TPC respectfully submits the following article-by-article technical analysis of the draft Implementing Regulation, focusing on the critical intersection between data governance and the evaluation of artificial intelligence models.

Article 2: Access to general-purpose AI models

Technical Critique:

- Overbroad Scope and Sequencing: The draft wording grants nearly unlimited technical penetration into a provider’s infrastructure without sufficient safeguards, thresholds, or proportional sequencing. It implies the Commission may immediately deploy the most intrusive forms of access rather than escalating only when lower-intrusion methods are inadequate.
- IP and Vulnerability Exposure: Assuming access implies "read-only" rather than direct "commit/edit" rights to source code or training data, deep-tier exposure still introduces severe intellectual property (IP) risks. Exposing the underlying architecture allows third parties to identify system prompts and logical loopholes, creating severe cybersecurity vulnerabilities that could be exploited later.
- Data Leakage: It remains ambiguous whether Article 2(1) implies direct access to training data. Even if excluded, granting comprehensive access to model weights and states makes the leakage of sensitive training data highly probable.
- Operational Stability: Conducting intrusive evaluations on live production infrastructure poses security risks; it also directly degrades the end-user experience. Live experimentation increases latency and query denial rates due to transient

adjustments to model parameters (e.g., altering thresholds for harmful content detection).

- Forensics and Lifecycle Gaps: Disabling logging completely removes the primary forensic defense for detecting unauthorized third-party intrusions. Furthermore, the draft omits critical temporal limits for how long access may be maintained and lacks provisions for the secure destruction of data once the evaluation concludes.

Proposed Refinements:

1. Proportional Escalation Protocol: Mandate a sequenced framework in which the Commission is encouraged to exhaust lower-intrusion evaluation methods (e.g., standard API queries) before demanding access to deep infrastructure, source code, or weights.
2. Containment & Scoping: Require the use of Technical Clean Rooms or Secure Mirror Environments for evaluations to protect live end-user environments. The regulation should explicitly clarify that any access to source code or data pipelines is strictly "read-only."
3. Data Lifecycle Boundaries: Enforce strict temporal limits on the duration of access and mandate the secure, verifiable destruction of any data or architectural insights obtained by the Commission or independent experts immediately upon the evaluation's conclusion.
4. Encrypted Shadow Logging: To balance confidentiality with security, permit Encrypted Shadow Logging. Access logs should be recorded, immediately encrypted, and placed in an immutable escrow, ensuring forensic auditability without compromising the immediate confidentiality of the Commission's specific queries.

Article 3: Independent experts

Technical Critique:

- Arbitrary Timelines and IP Seepage: A fixed 12-month retrospective independence check is structurally inadequate for mitigating the risks of intellectual property (IP) transfer. Experts granted deep-tier access to proprietary model weights, source code, and data curation techniques acquire durable, highly sensitive technical knowledge.
- Lack of Ongoing Verification: The draft mandates that an expert "shall remain independent" but provides no enforcement or verification mechanisms to ensure this holds true throughout the evaluation. The absence of continuous disclosure obligations leaves a critical vulnerability.
- Absence of Due Process: The current text specifies no procedural mechanism for a provider to formally challenge, object to, or request the recusal of an appointed expert if a latent conflict of interest or significant IP risk is identified.

Proposed Refinements:

1. Risk-Tiered Post-Evaluation Restrictions: Rather than relying on arbitrary temporal limits (e.g., 12 or 24 months), the Commission should implement a risk-tiered framework. Experts granted internal access to non-public model architectures and data pipelines should be bound by stringent legal and contractual assurances that explicitly prohibit consulting or employment agreements with competing general-purpose AI providers. The severity of these legal restrictions should scale in proportion to the depth of technical access granted.
2. Continuous Disclosure Obligations: Introduce mandatory, continuous verification mechanisms, including a formal mid-appointment self-declaration of independence, to guarantee that no conflicts of interest materialize during the evaluation lifecycle.
3. Formal Objection Mechanism: Establish a transparent procedural framework allowing providers to challenge an expert's appointment. Providers should have a secure avenue to present evidence of potential IP compromise or conflicts of interest before the expert is granted access to sensitive infrastructure.

Article 5: Opening of proceedings

Technical Critique:

- Systemic Cascade Failures: General-purpose AI models serve as foundational infrastructure. Abruptly suspending a model's availability without accounting for technical dependencies risks triggering systemic cascade failures, data desynchronization, and severe operational disruptions across the downstream digital ecosystem.
- Lack of Risk Stratification: The draft treats all grounds of urgency identically. It fails to distinguish between structural non-compliance (which requires a controlled transition) and immediate, catastrophic behavioral risks (e.g., adversarial exploitation or rapid-onset toxic generation) where controlled mitigation is too slow.
- Jurisdictional Fragmentation: There is no technical or legal clarity regarding how unilateral interim measures by the Commission interact with parallel proceedings conducted by national competent authorities in the EU or in other non-EU jurisdictions (e.g., the UK). Uncoordinated emergency suspensions risk fracturing cross-border data pipelines and creating conflicting compliance mandates.

Proposed Refinements:

1. Technical Dependency Assessment with "Severe Risk" Override: Mandate a rapid Technical Dependency Assessment before enforcing market suspension to ensure that technically viable transition protocols are established for downstream operations. However, the regulation should explicitly state that in cases of immediate and severe risk, instantaneous suspension takes absolute precedence over controlled transition measures.

2. Cross-Border Synchronization Protocols: Establish explicit procedural frameworks for coordinating interim measures with EU Member State national competent authorities and non-EU international regulatory bodies. This ensures that emergency model suspensions are synchronized, preventing operational chaos in parallel jurisdictional proceedings.

Article 7: Written observations on preliminary findings

Technical Critique:

- Inadequate Timeframe for Forensic Analysis: A 14-day window is technically inadequate for providers to conduct rigorous forensic investigations into general-purpose AI systems. While a two-week period may be sufficient for trivial administrative tasks or basic procedural checks, it fundamentally underestimates the complexity of algorithmic auditing.
- Resource Intensity of Technical Validation: Validating regulatory claims regarding multi-dimensional data bias, latent capabilities, or complex model failures requires extensive computational resource allocation, retraining simulations, and deep-tier statistical verification that cannot be responsibly executed within 14 days.

Proposed Refinements:

1. Tiered Response Framework: The regulation should implement a stratified timeline based on the technical severity of the findings. The 14-day minimum baseline should be strictly limited to administrative, procedural, or trivial observations. Conversely, technical findings that necessitate the re-evaluation of training data, model weights, or system architectures should carry a minimum response window of 45 to 60 days to ensure scientific accuracy and a rigorous defense.

Articles 8 & 9: Access to the file and Confidentiality

Technical Critique:

- Unnecessary IP Exposure: Granting completely unredacted access to raw technical files, architectural documents, and training logs to external experts poses significant risks of intellectual property exposure and potential data privacy violations.
- Divergence from Standard Audit Practices: The current draft diverges from established industry practices for deep-tier technical audits. Providing entire unredacted files bypasses fundamental "need-to-know" access controls, unnecessarily exposing core business secrets and algorithmic know-how when targeted access would suffice for the evaluation.

Proposed Refinements:

1. Industry-Standard Audit Precautions: The regulation should codify established operational safeguards routinely used by technology companies during high-stakes compliance and security audits. Specifically, document disclosure and access protocols should explicitly align with recognized international control frameworks, such as **ISO/IEC 27001** (Information Security Management), the **NIST Cybersecurity Framework (CSF)**, and **SOC 2** (System and Organization Controls) principles, to ensure confidentiality and strict access governance.
2. "Need-to-Know" Access Controls: Restrict unredacted file access strictly to the specific document sections or data logs absolutely necessary to evaluate the immediate compliance concern. Unrelated proprietary know-how embedded in the same files should remain protected.
3. Enhanced Legal Safeguards: Mandate that external legal, economic, and technical experts are bound by scientifically and legally rigorous reassurance clauses. This should include strict liability confidentiality agreements tailored to protect AI architectural know-how and prevent the lateral transfer of trade secrets.

Article 14: Transmission and receipt of information

Technical Critique:

- Ambiguous Terminology ("Methods"): The draft's mandate that the Commission shall define the "methods" for API sharing is vague. It fails to specify whether this refers to transmission protocols, data formats, or specific HTTP methods. Without explicitly restricting access to safe, read-only operations (e.g., HTTP GET), the current wording could theoretically allow the Commission to demand endpoints that permit state modification (e.g., POST/PUT/DELETE), which is unacceptable for compliance monitoring.
- Resource Contention and Overhead: The continuous, real-time extraction of data via regulatory APIs introduces severe resource contention risks. Unregulated automated querying by the Commission's evaluation systems effectively acts as an internal Denial of Service (DoS) vector. It consumes critical compute cycles and bandwidth that the AI model needs to perform its intended functions, degrading latency and availability for legitimate end users.

Proposed Refinements:

1. Technical Specificity for "Methods": The regulation should explicitly define "methods" as strictly "read-only transmission protocols." The Commission’s API access should be architecturally restricted to ensure it cannot inadvertently alter system states, modify configurations, or pollute live data pipelines.
2. Strict Rate Limiting and Query Caps: Rather than imposing additional overhead burdens on providers, the regulation should enforce strict limit-caps on the regulator. The Commission should commit to predefined Maximum Query Rates (e.g., a hard

cap of *X* requests per *Y* time period) and concurrency limits. This ensures that automated technical audits operate within safe, bounded resource limits, safeguarding the operational stability of the evaluated AI systems.

Acknowledgements:

ACM would like to thank the following members of the Europe TPC Committee for their contributions: Dr. Deniz Cetinkaya, Dr. Paolo Giudici, Dr. Gerhard Schimpf, Mr. Alejandro Saucedo, Dr. Eirini Ntoutsi, and the Chair of the AI Sub-Committee, Mr. Gaston Besanson. We would also like to thank the Europe TPC Committee for their review of these recommendations.