MOHAMMED ARBAAZ SHAREEF

Proposal and Comment Information

Title: Anti-Money Laundering and Countering the Financing of Terrorism Programs, R-1835

Comment ID: FR-0000-0133-02-C19

Submitter Information

Name: Mohammed Arbaaz Shareef

Submitted Date: 08/11/2026

I respectfully submit this comment in my personal capacity, with my full comment attached as a PDF. I am a data engineer with approximately a decade of experience building real-time data infrastructure in regulated U.S. financial environments, including Apache Kafka-based streaming platforms at a U.S. consumer lender operating more than 400 branches. I hold the CAMS designation and the Google Cloud Professional Data Engineer certification, and I am the sole inventor on three U.S. provisional patent applications directed to explainable, real-time, graph-based detection of financial crime (U.S. Provisional Applications No. 63/999,870; No. 64/001,569; and No. 64/001,564, filed March 2026). The views expressed are my own.

I write in support of the proposed rule's effectiveness framework and its treatment of innovative technology, and offer technical perspective on five of the Board's questions.

General support (Section IV.B): The proposal's recognition of machine learning, GenAI, digital identity, and analytics, and its assurance that banks responsibly incorporating such technologies will not incur additional supervisory or enforcement risk on that basis, remove a genuine barrier to adoption. In my experience, the principal reason advanced detection models are not deployed is not technical immaturity; it is uncertainty about whether they can survive validation, examination, and audit. I encourage the Board to retain this language and to pair it with a statement of what responsible incorporation looks like: models whose outputs are explainable to investigators, validators, and examiners, and whose decisions are preserved in auditable records.

Question 1: A monitoring model can be part of a reasonably designed program when its outputs can be explained (attribution of the transaction paths, network relationships, and features contributing to a risk score, with a plain-language rationale usable in SAR narratives) and when each scoring decision is preserved with sufficient metadata (model version, data snapshot, timestamp, case linkage) for after-the-fact review. Stating this would align the effectiveness framework with model governance.

Questions 15 and 16: In batch architectures, the feedback loop between risks and risk assessment is measured in weeks or months. Event-driven streaming architectures invert this: each new transaction, account change, or relationship change triggers re-scoring of the affected portion of the institution's customer and transaction graph, and analyst dispositions (for example, confirmed false positives) recalibrate detection continuously. I encourage the Board to draft the "promptly" update standard in a technology-neutral way that recognizes continuous, event-driven risk assessment as fully satisfying the requirement.

Question 17: Independent testing of ML-based monitoring is only meaningful if the model exposes why it scores what it scores. Testing should encompass review of the model's explanation artifacts - feature and relationship attributions, typology coverage against the risk assessment, and the auditable record of scoring decisions. The presence of such artifacts is evidence of a testable, and therefore maintainable, program.

Question 19: For model-based monitoring, the faithful written record of "transaction monitoring rules and parameters" is the machine-generated one: model version and configuration metadata, training-data lineage, explanation outputs attached to each alert or score, and immutable logs linking decisions to cases. I urge the Board to recognize such system-generated electronic records as satisfying written-documentation expectations.

Question 28: Smaller institutions need access to network-level detection, not exemption from it. The detection failures with the greatest national consequence are hardest to avoid at institutions with the least tuning, validation, and model-development capacity. Emerging privacy-preserving techniques - including federated approaches in which institutions collaboratively improve shared detection capability while raw

customer records never leave any institution and only encrypted model updates are exchanged - can give community banks and credit unions access to network-level detection no single small institution could build alone. I encourage the Board to confirm that participation in such privacy-preserving collaborative arrangements is consistent with a risk-based program and the responsible-innovation assurance in Section IV.B.

The full discussion of each point is in the attached PDF. I appreciate the Board's consideration of these comments.

Mohammed Arbaaz Shareef, CAMS
Submitted in my personal capacity.

Mohammed Arbaaz Shareef

August 11, 2026

Benjamin W. McDonough, Secretary
Board of Governors of the Federal Reserve System
20th Street and Constitution Avenue NW
Washington, DC 20551

Re: Docket No. R-1835; RIN 7100-AG78 — Anti-Money Laundering and Countering the Financing of Terrorism Programs (Notice of Proposed Rulemaking, 91 FR 42363, July 9, 2026)

Dear Secretary McDonough:

I respectfully submit this comment in my personal capacity. I am a data engineer with approximately a decade of experience building real-time data infrastructure in regulated U.S. financial environments, including streaming (Apache Kafka-based) platforms at a U.S. consumer lender operating more than 400 branches. I hold the Certified Anti-Money Laundering Specialist (CAMS) designation and the Google Cloud Professional Data Engineer certification, and I am the sole inventor on three U.S. provisional patent applications directed to explainable, real-time, graph-based detection of financial crime (U.S. Provisional Applications No. 63/999,870; No. 64/001,569; and No. 64/001,564, filed March 2026). The views expressed here are my own. I write to support the proposed rule's effectiveness framework and its treatment of innovative technology, and to offer technical perspective on several of the Board's specific questions.

I. General support: the effectiveness framework and the innovation language (Section IV.B)

The proposed rule's recognition that innovative approaches "could involve machine learning, generative artificial intelligence (GenAI), digital identity, blockchain monitoring and analytics, or application programming interfaces," and its assurance that banks responsibly incorporating such technologies "will not incur on that basis any additional risk" of significant supervisory or enforcement action, remove a genuine barrier to adoption. In my experience, the principal reason advanced detection models are not deployed is not technical immaturity; it is uncertainty about whether such models can survive validation, examination, and audit. I therefore encourage the Board to retain this language in the final rule and to pair it, in the preamble or subsequent guidance, with a statement of what "responsible" incorporation looks like: models whose outputs are explainable to investigators, validators, and examiners, and whose decisions are preserved in auditable records. Explainability is the property that reconciles the AML Act's innovation purpose (section 6002(3)) with the Board's long-standing model risk management expectations; making that connection explicit would accelerate responsible adoption more than any other single clarification.

II. Question 1 (effective AML/CFT program): clarify that explainability and auditability are the bridge between advanced models and "reasonably designed"

The description of an effective program is generally clear. One addition would materially help banks evaluating machine-learning detection: a statement that a monitoring model is capable of being part of a "reasonably designed" set of internal policies, procedures, and controls when its outputs can be explained — for example, through attribution of the specific transaction paths, network relationships, and features that contributed to a risk score, together with a plain-language rationale usable in SAR narratives — and when each scoring decision is preserved with sufficient metadata (model version, data snapshot, timestamp, case linkage) to permit after-the-fact review. Without such a statement, banks reasonably fear that stronger

but more complex models will be treated as unexaminable black boxes; with it, the effectiveness framework and model governance point in the same direction.

III. Questions 15 and 16 (prompt updates; the feedback loop between risks and risk assessment): event-driven architectures make "promptly" technically achievable

The Board asks whether the "promptly" update standard would change current practice and how risks and risk assessment processes feed one another. I offer a technical observation: in traditional batch architectures, detection scenarios run on overnight or longer cycles, and risk assessments are revisited on periodic schedules; the feedback loop the Board describes in Question 16 is measured in weeks or months. Modern streaming architectures invert this. When detection is triggered by data mutations — that is, when each new transaction, account change, or relationship change initiates re-scoring of the affected portion of the institution's customer and transaction graph — the institution's operative picture of its risk updates continuously, and analyst dispositions (for example, confirmed false positives) can recalibrate detection on an ongoing basis rather than at annual tuning exercises. I encourage the Board to draft the "promptly" standard in a technology-neutral way that recognizes continuous, event-driven risk assessment as fully satisfying — indeed exceeding — the requirement, so that institutions adopting real-time approaches are not asked to also maintain parallel calendar-based processes that add cost without adding effectiveness.

IV. Question 17 (independent testing): explainable outputs are what make ML-based monitoring testable

Independent testing of a scenario-based program is well understood: the tester reviews rule logic, thresholds, and coverage. Testing a machine-learning-based program is only meaningful if the model exposes why it scores what it scores. I encourage the Board, in the final rule or accompanying guidance, to note that for ML-based monitoring, independent testing should encompass review of the model's explanation artifacts — feature and relationship attributions, coverage of typologies against the institution's risk assessment, and the auditable record of scoring decisions — and that the presence of such artifacts is evidence of a testable, and therefore maintainable, program. This framing gives banks a concrete design target and gives testers objective criteria, consistent with the proposal's emphasis on testing focused on effectiveness.

V. Question 19 (written program): recognize machine-generated audit records as satisfying documentation expectations

The Board asks which components must be written and in what form, specifically including "transaction monitoring rules and parameters" and whether documentation may take the form of "system configurations[] or electronic records." I urge the Board to answer yes: for model-based monitoring, the faithful written record of "rules and parameters" is the machine-generated one — model version and configuration metadata, training-data lineage, the explanation outputs attached to each alert or score, and immutable logs linking decisions to cases. Narrative summaries describing a model in prose age quickly and diverge from the deployed system; electronic records generated by the system itself are contemporaneous, complete, and examiner-verifiable. Recognizing such records as "written" documentation would align the rule with how modern detection systems actually preserve their reasoning, and would encourage institutions to build auditability in from the start rather than reconstruct it for examinations.

VI. Question 28 (tailoring to size and complexity): smaller institutions need access to network-level detection, not exemption from it

The preamble candidly notes that adopting new technologies "may not be suitable for all banks, particularly smaller ones." I respectfully suggest the final rule frame this not only as a matter of proportionate

expectations but also as a capability-access problem the rule can help solve. The detection failures with the greatest national consequence — including monitoring systems that fail to capture material volumes or types of transactions, an example the preamble itself gives of material implementation failure — are hardest to avoid precisely at institutions with the least tuning, validation, and model-development capacity. Two features of the proposal help and deserve emphasis: first, the preamble's recognition that risk assessment processes may draw on information obtained from other financial institutions, including through section 314(b) sharing; second, the effectiveness framework's technology neutrality. Emerging privacy-preserving techniques — including federated approaches in which institutions collaboratively improve shared detection capability while raw customer records never leave any institution and only encrypted model updates are exchanged — can give community banks and credit unions access to network-level detection that no single small institution could build alone, within the confidentiality constraints that have prevented raw-data consortium approaches. I encourage the Board to confirm, in guidance or the preamble to the final rule, that participation in such privacy-preserving collaborative arrangements is consistent with a bank's risk-based program and with the responsible-innovation assurance in Section IV.B.

VII. Conclusion

The proposed rule's effectiveness orientation, its explicit welcome of machine learning and related technologies, and its assurance against innovation-based enforcement risk together advance the AML Act's purpose of encouraging technological innovation against money laundering and the financing of terrorism. My recommendations are offered to sharpen that direction: make explainability and auditability the stated hallmarks of responsible model adoption (Questions 1 and 17); draft the update standard to recognize continuous, event-driven risk assessment (Questions 15 and 16); treat machine-generated audit records as written documentation (Question 19); and affirm privacy-preserving collaboration as a path for smaller institutions (Question 28). I appreciate the Board's consideration of these comments.

Respectfully submitted,

Mohammed Arbaaz Shareef, CAMS
Data Engineer; sole inventor, U.S. Provisional Patent Applications No. 63/999,870, No. 64/001,569, and No. 64/001,564
Submitted in my personal capacity; the views expressed are my own.