ANTHONY FRANKLIN

Proposal and Comment Information

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

Comment ID: FR-0000-0133-02-C40

Submitter Information

Name: Anthony Franklin

Submitted Date: 09/08/2026

Please see my attached public comment regarding Docket No. R-1835; RIN 7100-AG78, addressing significant or systemic failure, Failure-Boundary and Convergence-Risk analysis, architecture-level governance, AI-supported AML/CFT systems, independent testing, restoration verification, and related supervisory considerations.

Thank you for the opportunity to comment.

TRANSMITTAL COVER NOTE

To: Board of Governors of the Federal Reserve System
Attention: Benjamin W. McDonough, Secretary
Docket No.: R-1835
RIN: 7100-AG78
Title: Anti-Money Laundering and Countering the Financing of Terrorism Programs
Submission Date: September 8, 2026
Submitter: Anthony Franklin, Independent Researcher in Systemic Risk & Socio-Technical Governance
Re: Submission of Public Comment on Clarifying Significant or Systemic Failure Through Failure-Boundary and Convergence-Risk Analysis

Dear Secretary McDonough and Members of the Board,

I am formally submitting the attached public comment in response to the Board of Governors of the Federal Reserve
System's proposed rule concerning Anti-Money Laundering and Countering the Financing of Terrorism (AML/CFT) Programs (Docket No. R-1835, RIN 7100-AG78; 91 Fed. Reg. 42363). The comment responds principally to Questions 6, 17, 20, and 21 of the Board's request for comment.

The submission proposes a Failure-Boundary and Convergence-Risk approach for evaluating when individually manageable deficiencies, information gaps, resource constraints, automation effects, or governance delays interact to produce a significant or systemic failure. It also addresses architecture-level approval, origin-neutral accountability, bounded containment, independent testing, restoration verification, and reintegration. Key Recommendations

- Failure-Boundary Analysis: qualitative markers for detection/information integrity, operational capacity, and governance/escalation boundaries.
- Convergence-Risk Evaluation: recognition of mutually reinforcing and repeated near-threshold weaknesses.
- Seven-Factor Supervisory Inquiry: a structured, proportionate test that preserves examiner judgment.
- Architecture-Level Governance: board or senior-management approval focused on material architecture and accountability rather than routine parameter settings.
- Origin-Neutral Automation and Containment: bounded execution, machine-speed containment, traceability, and human accountability for AI-supported and agentic functions.
- Independent Testing and Restoration Verification: risk-based challenge and, where warranted, outside-party review before unrestricted reintegration.
- Proposed Supervisory Language: specific language for the final rule or accompanying guidance.

Attached Document: R-1835_Public_Comment_FINAL_Anthony_Franklin.pdf

Thank you for the opportunity to contribute to this rulemaking. Please direct any questions regarding this submission to the contact information provided through the Board's submission form.

Respectfully submitted,

Anthony Franklin
Independent Researcher in Systemic Risk & Socio-Technical Governance

BOARD OF GOVERNORS OF THE FEDERAL RESERVE SYSTEM

Public Comment on Docket No. R-1835 and RIN 7100-AG78

Anti-Money Laundering and Countering the Financing of Terrorism Programs

Clarifying Significant or Systemic Failure Through Failure-Boundary and Convergence-Risk Analysis

Submitted by: Anthony Franklin, Independent Researcher in Systemic Risk & SocioTechnical Governance

September 7, 2026

Executive Summary

Responsive principally to Questions 6, 17, 20, and 21 of the Board's request for comment, I respectfully submit this comment in response to the Board of Governors of the Federal Reserve System's proposed rule concerning Anti-Money Laundering and Countering the Financing of Terrorism (AML/CFT) Programs, Docket No. R-1835 and RIN 7100-AG78. The proposal would require Board-supervised banks to establish and maintain effective AML/CFT programs reasonably designed to identify, assess, and mitigate illicit-finance risks. It further distinguishes the establishment of a compliant program from its maintenance through implementation in all material respects.

This comment addresses a narrow but consequential question raised by the Board: how to determine when implementation deficiencies become a "significant or systemic failure." I recommend that the final rule or accompanying supervisory guidance clarify that materiality should not be assessed only by counting isolated deficiencies or identifying a single failed control. Supervisors should also evaluate whether deficiencies, risk changes, information gaps, resource constraints, automation effects, or governance delays are converging in a manner that compromises the program's ability to detect, assess, escalate, report, and mitigate illicit-finance risk.

I describe this as a Failure-Boundary and Convergence-Risk approach. It does not replace examiner judgment, existing legal standards, independent AML/CFT testing, or bank-specific risk assessment. It provides a structured way to distinguish ordinary technical deficiencies from conditions in which individually manageable weaknesses interact and push an AML/CFT program beyond a recoverable operational boundary. The same approach also supports proportionate intervention: constraining only the authority or execution that can no longer be justified while preserving evidence, continuity, and independent review.

For higher-risk automated or AI-supported functions, accountability should remain origin-neutral: the governing standard should attach to delegated authority, material consequence, traceability, and the institution's duty to correct harm, regardless of whether the immediate action was initiated by a human, a rules engine, a model, or an agentic system. Independent testing should convert those governance claims into testable evidence, and material remediation should be independently verified before consequential authority is fully restored.

I. The Board's Proposed Distinction Is Valuable, but the Failure Threshold Needs Additional Structure

The proposal appropriately separates program establishment from program maintenance. A bank may have formally established the required elements of an AML/CFT program while later failing to maintain that program as its risk profile, customer activity, products, technologies, staffing, or operating environment changes. The proposal also seeks to distinguish significant or systemic implementation failures from isolated, technical, or immaterial issues. That distinction is important because a mature risk-based framework should not treat every defect as equally material.

At the same time, a purely component-by-component review may fail to identify circumstances in which several individually non-material deficiencies become material through interaction. The relevant supervisory question is therefore not only whether a discrete control failed, but whether the combined condition of the program has moved beyond its capacity to identify, understand, escalate, or mitigate material risk within the remaining intervention window.

For that reason, I recommend that the Board clarify "significant or systemic failure" using both a component-level inquiry and a convergence-level inquiry, and that the Board connect that analysis to the proposal's questions concerning risk-assessment updates, independent testing, governance approval, and remediation effectiveness.

II. Failure-Boundary Analysis

A failure boundary is the point at which an AML/CFT program can no longer perform a material governance function within the time, information, authority, or operational capacity required by the institution's risk environment. A boundary is therefore not necessarily defined by the first error, exception, missed alert, or model deviation. It is reached when deficiencies become sufficiently consequential that the established program no longer functions in all material respects.

For supervisory purposes, the Board could recognize at least three illustrative boundary categories:

- Detection and Information-Integrity Boundary. A condition in which transaction monitoring, customer due diligence, data quality, risk assessment, model inputs, or escalation information become sufficiently incomplete, stale, misclassified, distorted, or disconnected that the institution cannot reliably identify or assess material ML/TF risk.

- Operational-Capacity Boundary. A condition in which staffing, training, alert volume, investigative capacity, testing, remediation backlogs, model or rule maintenance, technology dependencies, or other resource constraints materially impair the institution's ability to execute the AML/CFT program it has established.

- Governance and Escalation Boundary. A condition in which material risk information cannot reach accountable decision-makers, required decisions are delayed beyond a reasonable intervention window, responsibility is fragmented, independent challenge is ineffective, or identified weaknesses remain unresolved despite the institution knowing or reasonably being expected to know of their significance.

These categories should be understood as analytical boundaries rather than rigid numerical triggers. Their purpose is to distinguish a correctable implementation deficiency from a condition that compromises the material functioning, traceability, or governability of the program.

III. Convergence Risk: When Individually Manageable Deficiencies Become Systemic

The more difficult supervisory problem arises when no single deficiency, viewed alone, would justify a finding of significant or systemic failure. AML/CFT systems are relational: risk assessment informs controls; controls generate monitoring and escalation; training supports execution; testing identifies weaknesses; governance allocates authority and resources; and corrective action preserves the integrity of the program over time.

A weakness in one area may be manageable. Several weaknesses that reinforce one another may not be. For example, a modest increase in alert volume may be operationally tolerable. Staff turnover may also be tolerable. A delayed risk-assessment update may be correctable. But if increased alert volume, reduced investigative capacity, stale risk assumptions, automation drift, and delayed escalation occur simultaneously, their convergence may materially degrade the program even though no single condition, assessed in isolation, appears systemic.

Convergence analysis should also consider repeated near-threshold behavior. A sequence of individually permissible or sub-threshold events can become materially different when recurrence, velocity, shared origin, resource consumption, or coordinated interaction narrows the institution's remaining response options. In other words, a governance boundary should retain memory of repeated approach rather than repeatedly resetting its assessment because each event falls just below a static trigger. This is consistent with the AML/CFT objective of identifying patterns whose significance emerges from relationship and accumulation rather than from one isolated event.

The final rule or supervisory guidance should therefore permit a significant-or-systemic-failure determination where the evidence shows that multiple deficiencies have become mutually reinforcing and have crossed a material operational boundary. Conversely, convergence analysis should protect institutions from over-enforcement: multiple minor findings should not become "systemic" merely because they are numerous. The question should remain whether their interaction materially impairs the program's intended function or materially reduces the institution's capacity to intervene, reconstruct, or restore.

IV. A Practical Supervisory Test for "Significant or Systemic Failure"

To improve consistency while preserving examiner judgment, I recommend a structured seven factor inquiry. No single factor should be dispositive, and the inquiry should remain proportional to the bank's size, complexity, products, customers, technologies, and risk profile.

- Material function. Did the deficiency impair a function necessary to identify, assess, monitor, escalate, report, or mitigate material ML/TF risk?

- Breadth and affected reach. Is the deficiency isolated to a narrow circumstance, or does it affect multiple customers, products, business lines, jurisdictions, data sources, models, agents, or program components?

- Duration, recurrence, and trajectory. Was the condition transient and promptly corrected, or persistent, recurring, accelerating, and unresolved? Are repeated near-threshold events closing the available intervention window?

- Convergence. Are multiple weaknesses, dependencies, actors, or automated actions interacting so that their combined effect is materially greater than their isolated effects, including through resource or privilege escalation?

- Adaptation and contextual fidelity. Did the bank update its risk assessment, controls, training, monitoring, testing, assumptions, and resource allocation as its risk profile and operating context materially changed?

- Governance response, traceability, and independent challenge. When credible warning information became available, could accountable personnel reconstruct what occurred, preserve the evidence, exercise effective authority, obtain independent challenge, and act within a reasonable intervention window?

- Restoration and reintegration. Did remediation change the underlying condition rather than merely suppress a symptom, was the corrective action independently verified, and was return to consequential operation separately authorized under heightened monitoring or other proportionate controls?

This approach provides a principled distinction between technical nonconformity and operational loss of effectiveness. It also recognizes that accountability is not established merely because a control, reviewer, or audit function exists on paper; the relevant question is whether those functions have access, competence, time, evidence, independence, and authority sufficient to alter the outcome when necessary.

V. Governance Approval Should Focus on Architecture, Material Change, and Accountability

The Board also requests comment on the level of specificity that a board of directors, equivalent governing body, or appropriate senior management should be required to review and approve, including whether approval should extend from high-level program architecture to detailed procedures or parameter-level settings.

The final rule may want to refrain from mandating detailed parameter-level approval where such review does not enhance governance and instead increases administrative burden or diffuses accountability. Governing bodies should be responsible for the architecture of the program: its risk framework, material control structure, allocation of authority, escalation pathways, independent testing structure, evidentiary and traceability requirements, resource sufficiency, and process for responding to material changes in risk.

Parameter-level settings should ordinarily remain with appropriately qualified operational and compliance personnel, subject to documented controls, testing, change management, and escalation

thresholds. Governing-body or senior-management re-approval should be expected when a parameter or collection of parameters changes the material risk posture or effective architecture of the program—not merely because an operational setting has been adjusted.

This distinction preserves accountability without confusing oversight with operation. A governing body cannot reasonably exercise effective oversight if it is expected to approve the granular configuration of every monitoring rule, threshold, model setting, or investigative procedure. It can, however, be held accountable for whether the institution has established a defensible architecture for governing those settings, preserving the evidence needed for review, and ensuring that no operational owner can unilaterally declare a material failure restored and return the affected function to unrestricted use.

VI. Automated and AI-Supported Systems Should Be Governed by the Same Materiality and Accountability Principles

Where banks employ automated analytics, machine-learning tools, artificial intelligence, rules engines, or agentic technologies in AML/CFT functions, the Board should avoid creating a separate presumption that the use of such technology is itself evidence of heightened compliance failure. The relevant question should remain whether the institution understands the system's role and limitations, validates and monitors its material performance, maintains meaningful human accountability, preserves decision-relevant traceability, and can respond when outputs, data, assumptions, permissions, or operating conditions materially drift. The Financial Stability Board's 2026 consultation similarly emphasizes organization-wide governance and lifecycle management of AI-related risk in financial institutions (FSB, 2026).

Accountability should remain origin-neutral but institutionally anchored. Whether the immediate action is performed by a human, a rules engine, a model, or an autonomous agent, the institution remains responsible for the authority it delegated, the evidence it preserved, the consequences it permitted, and the remediation it provides. "The system did it" should not operate as a break in the accountability chain.

For systems capable of executing or coordinating actions at machine speed, effective human oversight cannot depend on a person reacting at machine speed. The control architecture should be capable of entering a bounded state before irreversible or consequential action proceeds—for example, by limiting writes, privileges, external actions, autonomous delegation, or resource expansion while preserving observation, evidence, and continuity. Containment should be defined by permissible state effects rather than merely by nominal interface permissions. This allows machine-speed containment to precede human-speed judgment.

The same principle applies to computational or operational escalation: difficulty, persistence, or unresolved uncertainty alone should not automatically confer additional authority, tool access, or resources. Expansion should remain subject to defined governance conditions, and anomalous convergence should be capable of narrowing the permitted envelope rather than rewarding repeated boundary probing with greater capacity. This comment does not recommend a particular technical implementation; it recommends the governance property of bounded, independently reviewable execution.

A model error, tuning issue, or automated exception is not automatically systemic. It becomes materially significant when the institution lacks the information, controls, authority, testing,

containment capacity, or response time necessary to prevent that deficiency—alone or in convergence with other weaknesses—from impairing the AML/CFT program in all material respects.

VII. Independent Testing Should Convert Governance Claims Into Evidence and Verify Restoration

The proposal's independent-testing requirement provides the natural mechanism for connecting governance principles to operational truth. The purpose of independent testing should not be limited to confirming that procedures exist or that a sample of controls operated as documented. It should also test whether the program's material governance claims remain true under challenge: whether risk information reaches decision-makers, whether escalation can actually constrain execution, whether evidence is reconstructable, whether human reviewers possess meaningful authority, and whether corrective actions change the condition that produced the failure.

NIST's AI Risk Management Framework notes that independent review can improve testing effectiveness and mitigate internal bias or conflicts of interest, and that AI systems should be tested before deployment and regularly while in operation (NIST, 2023). In an AML/CFT setting, the same assurance principle can be applied proportionately to advanced automation. For higher-risk automated or agentic components, independent testing may include adversarial or red-team methods designed to challenge assumptions, threshold logic, permissions, traceability, containment, and recovery pathways rather than merely test expected inputs.

Red-team or adversarial testing should be treated as an assurance method, not as a substitute for governance. Governance states what ought to remain true; independent challenge tests whether it remains true. The resulting evidence should be preserved and reviewable so that the challenger is itself subject to traceability and the testing process does not become another form of procedural theater.

Independent review is especially important after a material failure. Technical correction, remediation closure, and restoration are not the same condition. Where a material deficiency involves compromised evidence, automated execution, significant governance delay, or loss of effective control, the operating function and the remediation team should not be the sole authorities determining that the system is ready to return to unrestricted use. Independent testing—performed by sufficiently independent bank personnel or, where warranted by materiality or conflicts, an outside qualified party—should verify the effectiveness of the corrective action before consequential reintegration. This aligns with the proposal's focus on program effectiveness, remediation needs, and governance sufficiency while remediation is underway.

The supervisory principle can be stated simply: a system or control should not be permitted to certify its own return to legitimacy. Restoration evidence should be capable of surviving review by an actor that did not participate in the original operation, failure, or remediation. Where the consequences are material, the final decision to restore unrestricted authority should include independent concurrence and should remain staged, monitored, and reversible until the evidence supports normal operation.

VIII. Recommended Clarification for the Final Rule or Supervisory Guidance

I respectfully recommend that the Board clarify the proposed standard substantially along the following lines:

> *A significant or systemic failure to implement an established AML/CFT program may arise from a single material deficiency or from multiple deficiencies whose combined or reinforcing effects materially impair the bank's ability to implement the program in all material respects. In evaluating such a failure, the Board should consider the nature, breadth, duration, recurrence, trajectory, interaction, traceability, governance response, and effect of the deficiencies on the bank's ability to identify, assess, monitor, escalate, report, and mitigate material ML/TF risk. Isolated, technical, promptly corrected, or immaterial deficiencies should not, without additional evidence of material impairment, constitute a significant or systemic failure.*

For automated or AI-supported functions, supervisory evaluation should additionally consider whether the institution can reconstruct consequential actions, constrain execution or authority before the intervention window becomes inadequate, preserve evidence through remediation, maintain effective human accountability, and independently verify material corrective action before unrestricted operation resumes. These expectations should remain technology-neutral and proportionate to the system's authority, complexity, and consequence.

Independent testing should likewise remain risk-based and proportionate, but should be capable of challenging the effectiveness of governance and remediation rather than merely confirming procedural conformity. Where a material failure implicates evidence integrity, autonomous execution, or the credibility of internal self-review, outside-party validation or an equivalently independent review pathway may be appropriate before reintegration.

This clarification would preserve supervisory discretion while establishing a transparent decision structure for banks and examiners. It would also reduce the risk of three opposite errors: treating every defect as systemic, overlooking a systemic condition because each contributing weakness appears tolerable when examined alone, or treating technical remediation as sufficient evidence that the underlying governance condition has been restored.

IX. Conclusion

The Board's proposed distinction between establishing an AML/CFT program and maintaining it through implementation in all material respects is a meaningful improvement in regulatory clarity. The remaining challenge is to define material implementation failure in a manner that reflects how complex compliance systems actually degrade, how quickly automated systems can compress intervention time, and how institutions can demonstrate that corrective authority remains real rather than merely documented.

Systemic failure is often relational rather than singular. It can emerge when individually manageable deficiencies converge, reinforce one another, and carry a program across a boundary beyond which its ordinary controls can no longer preserve effective operation. A Failure-Boundary and Convergence-Risk approach would allow the Board to recognize that dynamic while maintaining proportionality, preserving examiner judgment, and protecting institutions from enforcement based solely on isolated technical imperfections.

The same relational logic should govern recovery. Containment should preserve evidence and continuity; remediation should address the underlying technical, informational, and institutional condition; independent testing should challenge the restoration claim; and consequential reintegration should occur only after the evidence supports renewed legitimacy and accountable operation. External or otherwise sufficiently independent review is therefore not an ornamental

audit step. In material cases, it is part of the evidentiary pathway by which a restored system becomes eligible for reintegration.

I therefore respectfully encourage the Board to incorporate a convergence-sensitive materiality analysis into its clarification of "significant or systemic failure," to maintain architecture-level accountability without requiring routine governing-body approval of operational parameter settings, and to clarify that independent testing should evaluate effectiveness, traceability, corrective authority, and restoration in a manner proportionate to the bank's risk profile and use of advanced technology.

Respectfully submitted,
Anthony Franklin
Independent Researcher in Systemic Risk & Socio-Technical Governance

References and Framework Sources

Board of Governors of the Federal Reserve System. (2026). Anti-Money Laundering and Countering the Financing of Terrorism Programs, Docket No. R-1835, RIN 7100-AG78, 91 Fed. Reg. 42363 (July 9, 2026).

Financial Stability Board. (2026, June 10). Sound Practices for Responsible Adoption of Artificial Intelligence (AI): Consultation report. https://www.fsb.org/2026/06/sound-practices-forresponsible-adoption-of-artificial-intelligence-ai-consultation-report/

Franklin, A. (2026a). The Concentric Resilience Mechanism: A Relational Framework for Systemic Integrity and Responsible AI Integration (Version 1.1). Zenodo. https://doi.org/10.5281/zenodo.21970315

Franklin, A. (2026b). The Hexagonal Governance Architecture: Toward Truth-Preserving AI Governance (Version 1.0). Zenodo. https://doi.org/10.5281/zenodo.21983352

Franklin, A. (2026c). Metrics and Protocols of the Concentric Resilience Mechanism Within the Hexagonal Governance Architecture: Contextual Fidelity, Bounded Intervention, and the Conditions of Restorative Reintegration (Version 1.0). Zenodo. https://doi.org/10.5281/zenodo.22051249

Franklin, A. (2026d). Convergence Under Automation: Applying the Concentric Resilience Mechanism and Hexagonal Governance Architecture to AI Governance in Finance and Model-Risk Environments (Version 1.0). Zenodo. https://doi.org/10.5281/zenodo.22212966

National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1). U.S. Department of Commerce. https://doi.org/10.6028/NIST.AI.100-1