KILLCHAIN, INC., D/B/A FLINT NETWORK, JOSEPH TAYLOR

Proposal and Comment Information

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

Comment ID: FR-0000-0133-02-C16

Submitter Information

Organization Name: KillChain, Inc., d/b/a FLINT Network

Organization Type: Company

Name: Joseph Taylor

Submitted Date: 08/04/2026

KillChain, Inc., doing business as FLINT, respectfully submits this comment, with a full comment letter attached, on the Board's Notice of Proposed Rulemaking, "Anti-Money Laundering and Countering the Financing of Terrorism Programs," 91 FR 42363 (July 9, 2026), Docket No. R-1835, RIN 7100-AG78. FLINT filed substantially similar comments with FinCEN (Docket No. FINCEN-2026-0034) and with the OCC, FDIC, and NCUA on the concurrent April 10, 2026 proposals, and has separately commented on the Board's joint stablecoin customer identification program proposal (Docket No. R-1885).

FLINT supports the proposed rule. The shift from process-based to effectiveness-based program requirements is the most consequential modernization of BSA/AML program rules in roughly 25 years, and it is the correct one. We particularly support the proposal's two-part effectiveness construction, under which a program must be both properly established and implemented in all material respects, and the Board's express statement that responsible adoption of innovative technologies, including machine learning, digital identity, blockchain monitoring and analytics, and APIs, will not by itself draw additional supervisory attention.

We write to identify one gap the final rule and its examination procedures should address. Autonomous AI agents are being granted delegated spending authority across card, ACH, and stablecoin rails on rails that Board-supervised banks operate. The existing compliance stack was built to verify humans. Know Your Customer establishes who the customer is; nothing in that stack establishes what the non-human actor executing a transaction is, whether its principal actually authorized it, whether the transaction falls within the scope of its delegated authority, or whether the agent has been hijacked or repurposed after its credentials were issued. A bank can fully satisfy customer-level diligence and still have no answer to the question an examiner applying an effectiveness standard will eventually ask: how do you demonstrate effective monitoring of transactions your customers did not personally initiate? This is a measurement problem, and effectiveness is a measurement standard. A program cannot be shown to be implemented in all material respects with respect to a category of activity its controls were never designed to see.

The attached letter describes the primitive that closes the gap: Know Your Agent (KYA) verification, the non-human complement to KYC, performed in real time before a transaction executes, producing a cryptographically signed, retained verification record stating what was checked, what was observed, and what was decided. The record, not a bare risk score, is the compliance artifact: portable, tamper-evident, and intelligible to an analyst adjudicating an alert and to an examiner years later. This capability sits squarely inside the innovation categories the proposal already names.

FLINT recommends that the final rule, its guidance, and its examination procedures: (1) confirm that favorable treatment of responsible innovation extends to transaction-level verification of AI agents, including agent identity, delegated-authority, and behavioral verification; (2) recognize signed, retained, transaction-level verification records as documentary evidence of effectiveness for agent-initiated transactions, including in subsequent FFIEC examination manual updates; (3) direct that risk assessments treat transactions initiated by non-human actors with delegated financial authority as a distinct risk category, and clarify that the introduction of agent-initiated activity warrants review of the risk assessment; (4) coordinate with FinCEN, the OCC, the FDIC, the NCUA, and NIST so that supervisory expectations for agent verification remain aligned across agencies and with emerging federal technical standards, consistent with the digital identity questions posed in Docket No. R-1885; and (5) preserve technology-neutral and rail-neutral framing across card, ACH, and stablecoin rails, consistent with the perimeter established by the GENIUS Act.

The full argument, including responses to the Board's specific requests for comment on risk assessment updates, the AML/CFT Priorities, and training, is set out in the attached letter. FLINT would welcome the opportunity to discuss agent verification and the verification record standard with Board staff.

KillChain, Inc. (d/b/a FLINT)
flint.network | contact@flint.network

August 3, 2026

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

Submitted via the Board's proposal comment portal (federalreserve.gov/apps/proposals) Docket No. R-1835, RIN 7100-AG78

Re: Comment on the Notice of Proposed Rulemaking, "Anti-Money Laundering and Countering the Financing of Terrorism Programs," 91 FR 42363 (July 9, 2026), amending 12 CFR part 208 (Regulation H). Board Docket No. R-1835, RIN 7100-AG78.

To the Board of Governors of the Federal Reserve System:

KillChain, Inc., doing business as FLINT, respectfully submits this comment on the Board's proposed rule requiring Board-supervised banks to establish and maintain effectiveness-based anti-money laundering and countering the financing of terrorism (AML/CFT) programs, in alignment with the rules proposed on April 10, 2026 by the Financial Crimes Enforcement Network (FinCEN), the Office of the Comptroller of the Currency, the Federal Deposit Insurance Corporation, and the National Credit Union Administration. FLINT supports the proposed rule. FLINT filed substantially similar comments with FinCEN on the concurrent rulemaking (Docket No. FINCEN-2026-0034) and with the OCC, FDIC, and NCUA on their parallel proposals, and has separately commented on the Board's joint proposal establishing customer identification program requirements for permitted payment stablecoin issuers (Docket No. R-1885), a rulemaking directly related to the argument below. We write to renew, for the Board's docket, the one gap the final rule and its examination procedures should address, and to describe the transaction-level verification primitive that closes it: Know Your Agent (KYA) verification with signed, retained verification records for transactions initiated by autonomous AI agents. We also respond to several of the Board's specific requests for comment.

I. Interest of the Commenter

FLINT builds trust infrastructure for AI agents that move money: the rail-agnostic identity, authority, and evidence layer for autonomous economic actors. FLINT verifies an agent's financial authority before a transaction executes, detects behavioral anomalies, and emits a cryptographically signed verification record that the counterparty retains as evidence. These records are designed from the outset to function as examiner-ready compliance evidence.

The perspective offered here is grounded in direct operational experience with the Bank Secrecy Act and with the federal systems that implement it. FLINT's founder spent his career inside the U.S. national security, counter threat finance, and digital identity missions: as a Special Operations Officer, an ODNI Intelligence Officer, a U.S. Secret Service Special Agent, and a GS-15 Senior Technology Advisor at GSA Login.gov, the federal government's shared digital identity service. That career included working with BSA data and FinCEN's BSA E-Filing ecosystem as both a producer and a consumer of the reporting this rule governs. We know what makes a suspicious activity report useful to an investigator, and we know what makes compliance evidence persuasive to an examiner. We built FLINT to generate exactly that kind of evidence for a class of transactions the current framework does not yet name.

II. FLINT Supports the Effectiveness-Based Standard and the Board's Innovation Language

The shift from process-based to effectiveness-based program requirements is the most consequential modernization of BSA/AML program rules in roughly 25 years, and it is the correct one. Static, checklist-driven programs reward documentation of process rather than detection of illicit finance. An effectiveness standard, harmonized across the five concurrent proposals, gives Board-supervised banks room to adopt controls that actually work and gives the Board's examiners a basis to demand results rather than rituals.

Two features of the Board's proposal deserve particular support. First, the proposal's two-part effectiveness construction, under which a program must be both established in conformity with the rule and maintained through implementation in all material respects, correctly separates design failure from operational failure. That distinction only functions, however, if effectiveness can be measured, because "implemented in all material respects" is an evidentiary claim. Second, the Board expressly names the technologies through which responsible innovation is expected to arrive: machine learning, generative artificial intelligence, digital identity, blockchain monitoring and analytics, and application programming interfaces, and it commits that banks adopting such tools responsibly will not face additional supervisory attention solely on that basis. That commitment matters. It removes the supervisory ambiguity that has historically kept institutions on legacy tooling, and, as explained below, the named categories of digital identity and APIs are precisely where verification of non-human actors will be built.

We also support the concurrent proposals' preservation of human oversight at critical adjudication points, including OFAC determinations and SAR decisions, and the Board's requirement that the AML/CFT officer be located in the United States and accessible to the Board and FinCEN. Automation should expand what a human reviewer can see and prove; it should not displace the reviewer's judgment. Signed transaction-level verification records are what make that human oversight meaningful at machine speed, and what make it examinable after the fact.

III. The Gap: Effectiveness Cannot Be Demonstrated for Agent-Initiated Transactions Without a Transaction-Level Verification Primitive

A structural change in the payments landscape is underway during this comment period. Autonomous AI agents are being granted delegated spending authority across card, ACH, and stablecoin rails. Major payment networks, processors, and platform vendors shipped production agent-payment protocols in 2025 and 2026, and the GENIUS Act has brought compliant stablecoin rails squarely into the AML/CFT perimeter these concurrent rules modernize. The exposure is not distributed evenly: it will land first and hardest on the largest Board-supervised banking organizations, which are furthest along in shipping agent-payment integrations, whose customers are delegating payment authority to agents today, on rails those institutions operate.

The existing compliance stack was built to verify humans. Know Your Customer establishes who the human or the legal entity is. Nothing in that stack establishes what the non-human actor executing the transaction is, whether it was actually authorized by its principal, whether the transaction falls within the scope of its delegated authority, or whether the agent is still behaving consistently with its mandate or has been hijacked, impersonated, or repurposed after its credentials were issued. A Board-supervised bank can fully satisfy customer-level diligence and still have no answer to the question an examiner applying an effectiveness standard will eventually ask: how do you demonstrate effective monitoring of transactions your customers did not personally initiate?

This is a measurement problem, and effectiveness is a measurement standard.

A bank cannot demonstrate that its program is "implemented in all material respects" with respect to a category of activity its controls were never designed to see. For agent-initiated transactions, there is today no commonly recognized unit of evidence at the transaction level. That is the missing primitive, and the Board's examiners will feel its absence before most of the institutions they supervise do.

IV. The Missing Primitive: KYA Verification and the Signed Verification Record

Know Your Agent (KYA) is to non-human financial actors what KYC is to humans and legal entities, and it does not replace KYC; it completes it. Applied in real time, before a transaction executes, KYA evaluates whether an agent is known, authorized, in scope, operating from an acceptable environment, and behaving consistently with its mandate. FLINT's implementation evaluates 6 verification layers:

- **Principal Identity:** the human or entity on whose behalf the agent acts.
- **Agent Identity:** the verifiable identity of the agent itself.
- **Wallet Provenance:** the history and standing of the funding instrument.
- **Authorization Scope:** whether this transaction falls within the authority the principal actually delegated.
- **Environment Identity:** whether the agent is operating from an expected, acceptable environment.
- **Cross-Merchant Reputation:** whether the agent's observed behavior across the network is consistent with its mandate.

The output is not a bare risk score. It is a transaction decision, allow, step-up, review, or block, plus a cryptographically signed, retained verification record stating what was checked, what was observed, and what was decided at the moment of execution. The record is the compliance artifact. It is portable, independently verifiable, tamper-evident, and designed to support review, dispute, SAR narrative, and examination workflows. Under the cross-domain agent passport (CDAP) model, agent credentials issued in one platform's trust domain can be verified in another, so the evidentiary record does not fragment along vendor or rail boundaries.

This primitive sits squarely inside the innovation categories the Board's proposal already names. Agent identity verification is digital identity applied to non-human actors. Delegated-authority verification is performed over APIs at transaction time. Wallet provenance draws on blockchain monitoring and analytics. Behavioral-consistency monitoring is machine learning applied to a class of actor that produces far richer telemetry than a human customer ever could. The Board does not need to add a technology to its list to cover agent verification; it needs only to confirm that the list reaches the non-human actors now appearing on the rails its supervised banks operate.

The primitive also directly serves the human-oversight requirements of the concurrent proposals. A human analyst adjudicating an alert or making a SAR determination about agent activity today has almost nothing to review, and an examiner reviewing that analyst's work has less. A signed verification record gives both the complete decision context: which agent, on whose authority, within what scope, from what environment, behaving how, decided how, and why. Human-in-the-loop oversight of machine-speed commerce is only meaningful if the machine leaves evidence a human can read and an examiner can verify. That is precisely what the verification record is for.

V. Responses to the Board's Specific Requests for Comment

On risk assessment updates. The Board asks whether it should further clarify when a supervised bank must review or update its risk assessment processes.

Yes, in one respect: the Board should clarify, in the final rule or its accompanying guidance, that a customer's delegation of payment or spending authority to autonomous software agents, and a bank's own deployment of products or channels that accept agent-initiated instructions, are changes to the bank's ML/TF risk profile that warrant review of the risk assessment. Agent-initiated activity differs from the underlying customer's own activity in velocity, in attack surface (credential theft, agent hijacking, prompt-level manipulation), and in the evidence it leaves behind. Treating it as an undifferentiated extension of the customer is how it will be missed.

On guidance regarding the AML/CFT Priorities. To the extent the Board issues guidance on incorporating the AML/CFT Priorities into risk assessments, that guidance should acknowledge that several standing priorities, including fraud, cybercrime, and proliferation financing through virtual assets, now have an agentic expression, and that assessing them requires visibility into non-human actors.

On training. The proposed training requirement is sound. The Board should confirm that risk-based training may, for institutions exposed to agent-initiated activity, include recognition of agent-mediated transactions and the review of agent verification evidence, so that the analysts exercising the human oversight described above know what they are looking at.

On the AML/CFT officer and program accessibility. We support the U.S.-location and accessibility requirements as proposed and offer no further clarification, except to note that a program whose evidence is generated as signed, portable, machine-verifiable records is structurally easier to make "available upon request" than one whose evidence must be reconstructed from operational systems after the fact.

VI. Recommendations

FLINT respectfully recommends that the final rule, its accompanying guidance, and the examination procedures that implement it:

1. Confirm that the proposal's favorable treatment of responsible innovation, including the named categories of machine learning, digital identity, blockchain monitoring and analytics, and APIs, extends to transaction-level verification of AI agents: agent identity verification, delegated-authority verification, and behavioral monitoring of non-human actors.
2. Recognize signed, retained, transaction-level verification records as documentary evidence that a program is implemented in all material respects with respect to agent-initiated transactions, in the same way transaction monitoring records evidence effectiveness for human-initiated activity, including in the examination procedures and FFIEC BSA/AML Examination Manual updates that will follow the final rule.
3. Direct that risk assessment processes under the final rule account for transactions initiated by non-human actors operating with delegated financial authority as a distinct risk category, and clarify that the introduction of agent-initiated activity, by a customer or by the bank's own product surface, warrants review of the risk assessment.
4. Coordinate with FinCEN, the OCC, the FDIC, and the NCUA on the concurrent program rules, and with the National Institute of Standards and Technology, including the NIST Center for AI Standards and Innovation agent-standards initiative and the NCCoE work on software and AI agent identity and authorization, so that supervisory expectations for agent verification align across agencies and with emerging federal technical standards rather than diverging from them. Treatment here should also remain consistent with the digital identity and verifiable-credential questions the Board and FinCEN have posed in the PPSI customer identification program rulemaking (Docket No. R-1885), where the same verification primitive operates at the account boundary.

5. Preserve technology-neutral and rail-neutral framing, so that the same evidentiary expectations apply to agent-initiated activity across card, ACH, and stablecoin rails, consistent with the perimeter established by the GENIUS Act.

VII. Recommendations

The effectiveness standard arrives at the same moment autonomous agents begin moving money at scale, and the Board has already written the vocabulary that covers the answer: digital identity, APIs, machine learning, and blockchain analytics, adopted responsibly, without supervisory penalty. What remains is to name the risk those tools must be aimed at. If the final rule and its examination procedures recognize agent-initiated transactions as a risk to be measured and signed transaction-level verification records as the evidence that measures it, Board-supervised banks will adopt that control because it is the cheapest honest way to demonstrate effectiveness, and the United States will have set the evidentiary baseline for agentic finance before the rest of the world does.

FLINT appreciates the Board's work on this rulemaking and would welcome the opportunity to discuss agent verification, the verification record standard, or reference implementations with Board staff.

Respectfully submitted,

Joseph Taylor
Founder and Chief Executive Officer
KillChain, Inc. (d/b/a FLINT)
contact@flint.network | flint.network