Public Comment on FDIC Notice of Proposed Rulemaking — GENIUS Act Implementation

Submitted by: Empty Set LLC (Nevada Series LLC) **Contact:** Brice Love, Co-Founder **Email:** brice@emptysetllc.com **Date:** March 20, 2026 **Re:** Proposed Rule — Approval Requirements for Issuance of Payment Stablecoins by Subsidiaries of FDIC-Supervised Insured Depository Institutions, 90 FR 59409 (December 19, 2025), as extended by 91 FR 6138 (February 11, 2026), RIN 3064-AG20 **Submitted via:** Email to comments@fdic.gov

---

Executive Summary

Empty Set LLC respectfully submits this comment in response to the Federal Deposit Insurance Corporation's Notice of Proposed Rulemaking implementing the Guiding and Establishing National Innovation for U.S. Stablecoins Act ("GENIUS Act") for state-chartered insured depository institutions. We write to address the behavioral safety of AI systems that FDIC-supervised banks will deploy to comply with the Act's operational requirements — and to propose that the FDIC incorporate third-party behavioral safety credentialing into its implementing regulations.

The FDIC occupies a unique position among the federal banking agencies implementing the GENIUS Act. It supervises the most diverse population of institutions — from community banks with a single branch to large regional institutions with complex technology stacks — and it does so in coordination with fifty state banking regulators whose own AI governance expectations are evolving at different speeds. The FDIC's implementing regulations will set the AI governance floor not only for the institutions it directly supervises but, through the state examination coordination process, for the broader state-chartered banking system.

This comment addresses three concerns specific to the FDIC's supervisory context:

First, the **dual supervision challenge**. State-chartered banks operate under both FDIC and state banking department supervision. AI governance standards for stablecoin activities must be coordinated to avoid conflicting or duplicative requirements — particularly as states like Colorado, New York, and Connecticut are developing their own AI governance frameworks for financial institutions.

Second, the **deposit insurance dimension**. The FDIC's fundamental mission is protecting depositors. AI behavioral safety failures in stablecoin operations — discriminatory transaction blocking, reserve management errors, compliance failures that trigger enforcement actions — directly threaten the safety and soundness of insured institutions and, by extension, the Deposit Insurance Fund.

Third, the **community bank proportionality problem**. Like credit unions, many FDIC-supervised community banks lack the internal model risk management capacity to independently validate AI systems used for stablecoin compliance. Third-party behavioral safety credentialing provides a proportionate solution.

Empty Set LLC operates a standards development initiative focused on behavioral safety credentialing for autonomous AI agents, with emphasis on the insurance and financial services sectors. Our work addresses the gap between the Treasury's February 2026 Financial Services AI Risk Management

Framework and the practical tools that FDIC examiners need to assess AI governance at supervised institutions.

I. The AI Compliance Imperative

The GENIUS Act's operational requirements — continuous BSA/AML transaction monitoring, real-time reserve attestation, sanctions screening at blockchain velocity, consumer disclosure management — demand AI-driven automation. This is equally true for state-chartered banks as it is for national banks. The technology is the same. The compliance obligations are functionally identical. The AI behavioral safety risks are the same.

What differs is the supervisory context.

A. BSA/AML and the Speed of Stablecoin Transactions

Stablecoin transactions settle in seconds or minutes, not days. BSA/AML compliance for stablecoin activities requires AI systems that can evaluate transactions against OFAC sanctions lists, detect structuring and layering patterns, assess beneficial ownership, and generate SAR referrals — all at speeds incompatible with human-primary review processes.

These AI systems will make consequential compliance decisions autonomously. A transaction flagged as suspicious may be automatically delayed or blocked. A pattern identified as structuring may trigger an automatic SAR filing that, once submitted to FinCEN, cannot be withdrawn. These are irrevocable actions with immediate consequences for customers and for the bank's regulatory standing.

The behavioral safety of these AI systems — their accuracy, their bias characteristics, their robustness under novel market conditions, their failure modes under stress — determines whether GENIUS Act compliance is substantive or performative.

B. Reserve Management Under Stress

The GENIUS Act's reserve requirements demand that stablecoin issuers maintain reserves equal to or exceeding outstanding stablecoin supply, composed of eligible assets (Treasury securities, insured deposits, repurchase agreements). For state-chartered banks, reserve management during market stress events will require AI systems that make rapid decisions about asset liquidation, redemption processing, and regulatory communication.

An AI reserve management system that behaves differently under stress than under normal conditions — for example, liquidating reserves in a sequence that maximizes short-term liquidity but violates concentration limits, or halting redemptions based on market conditions that do not meet the statutory threshold for suspension — creates safety-and-soundness risk. Behavioral safety credentialing evaluates AI system performance under stress scenarios specifically, testing whether the system's decisions remain within defined operational boundaries when conditions deteriorate.

II. The Dual Supervision Challenge

State-chartered banks operate under a supervisory framework where the FDIC and the state banking department share examination authority. For stablecoin activities, this dual supervision creates coordination challenges that the GENIUS Act implementing regulations should address.

A. State AI Governance Frameworks Are Diverging

Several states have enacted or are developing AI governance requirements for financial institutions that are more prescriptive than current federal standards:

- **Colorado SB 24-205** (effective June 30, 2026, as delayed by SB 25B-004) requires deployers of high-risk AI systems to implement risk management policies, conduct impact assessments, and provide notice to affected consumers. Banks using AI for stablecoin compliance in Colorado will need to comply with these requirements in addition to federal standards once operative.
- **New York DFS Circular 7** requires NYDFS-supervised institutions to maintain board-level oversight of AI risk and to manage third-party AI vendor risk through documented governance programs.
- **Connecticut Bulletin MC-25** requires annual certification of AI governance practices for insurers — a model that state banking regulators may adopt for banks.

The FDIC's implementing regulations should establish a behavioral safety credentialing framework that serves as a common standard across both federal and state examinations. If the FDIC recognizes specific credentialing criteria, state banking departments can align their own AI governance expectations accordingly — reducing the compliance burden on state-chartered banks that would otherwise face inconsistent standards from their federal and state supervisors.

B. Examination Coordination

FDIC examiners and state examiners frequently conduct joint or alternating examinations of state-chartered banks. AI governance assessment for stablecoin activities should be conducted using consistent evaluation criteria so that a bank does not receive conflicting conclusions from its federal and state supervisors about the adequacy of its AI governance program.

Third-party behavioral safety credentials provide a common reference point. Both FDIC and state examiners can evaluate whether the bank's AI vendors hold current credentials, whether those credentials address the relevant behavioral safety dimensions, and whether the bank has appropriate procedures for monitoring credential status and responding to credential revocations. This is more practical than expecting both examination teams to independently evaluate proprietary AI systems whose architectures may be opaque to both.

III. The Deposit Insurance Dimension

The FDIC's unique mandate — protecting depositors and the Deposit Insurance Fund — adds a dimension to AI behavioral safety that other regulators do not share.

A. Safety and Soundness as AI Governance

AI behavioral safety failures in stablecoin operations create direct safety-and-soundness risk:

- **BSA/AML compliance failures** can trigger enforcement actions, civil money penalties, and in severe cases, cease-and-desist orders or consent decrees. An AI system that systematically underdetects suspicious activity exposes the bank to regulatory liability that threatens its viability.
- **Reserve management errors** can create liquidity crises. An AI system that makes flawed liquidation decisions during a stress event can amplify rather than mitigate redemption pressure, potentially threatening the institution's solvency.
- **Discriminatory false positives** can trigger fair lending scrutiny, Community Reinvestment Act concerns, and reputational damage that erodes the deposit base.

Each of these scenarios represents a threat to the Deposit Insurance Fund. The FDIC's implementing regulations should treat AI behavioral safety in stablecoin operations as a safety-and-soundness issue — not merely a technology management issue — and should incorporate AI governance assessment into the CAMELS framework for institutions engaged in stablecoin activities.

B. The Insurance Market Signal

The insurance industry's assessment of AI risk provides an independent signal that the FDIC should incorporate into its supervisory approach. The January 2026 ISO endorsements (CG 40 47, CG 40 48, CG 35 08) have restructured AI liability coverage across the commercial insurance market. Banks deploying AI systems for stablecoin operations face the same coverage contraction that is affecting all industries.

Institutions with demonstrable AI governance programs — including behavioral safety credentialing — will have access to affirmative AI liability coverage from the emerging specialty market (Lloyd's-backed MGAs such as Armilla AI and Testudo). Institutions without such programs will bear AI-related operational liabilities on their balance sheets without insurance risk transfer.

The FDIC should recognize that a bank's AI insurance coverage status — whether its AI-related operational liabilities are insured, self-insured, or uninsured — is a relevant safety-and-soundness indicator. Banks that have invested in behavioral safety credentialing and obtained affirmative AI liability coverage demonstrate a level of AI governance maturity that should be reflected positively in examination assessments.

IV. The Community Bank Proportionality Problem

The FDIC supervises thousands of community banks with total assets under $1 billion. These institutions lack the internal model risk management capacity to independently validate AI systems — but the GENIUS Act's compliance requirements do not scale down. A community bank's stablecoin transactions require the same BSA/AML monitoring, the same reserve management, and the same consumer protections as those of a large regional bank.

Third-party behavioral safety credentialing resolves this proportionality problem in the same way that SOC 2 attestations resolved the information security proportionality problem. The credentialing standard is applied to the AI system, not to the institution deploying it. A community bank using a credentialed AI system receives the same behavioral safety assurance as a large bank — without needing to maintain an internal team capable of replicating the evaluation.

The FDIC should explicitly recognize behavioral safety credentialing as a proportionate compliance

mechanism for community banks, ensuring that the AI governance expectations for stablecoin activities do not create a de facto barrier to entry that limits stablecoin services to only the largest and most technologically sophisticated institutions.

V. The Treasury FS AI RMF and the Validation Gap

The Treasury Financial Services AI Risk Management Framework, published February 19, 2026 by the Artificial Intelligence Executive Oversight Group (AIEOG) in coordination with the Financial Services Sector Coordinating Council (FSSCC) and the Financial and Banking Information Infrastructure Committee (FB-IIC), establishes 230 control objectives for AI governance in financial services. The framework adapts the NIST AI Risk Management Framework to the regulatory, operational, and fiduciary context of financial services.

The FS AI RMF establishes what institutions should do. It does not establish how the FDIC verifies compliance. FDIC examiners reviewing a state-chartered bank's stablecoin AI systems will need to assess compliance with control objectives that span governance, data management, model development, validation, monitoring, third-party risk, and consumer protection. Many of these objectives involve technical evaluations that require expertise in machine learning, statistical modeling, and adversarial testing.

Third-party behavioral safety credentialing bridges this validation gap. An independent credentialing body evaluates the AI system against the FS AI RMF's control objectives (or a subset relevant to the system's function) and issues a verifiable credential. The FDIC examiner can then focus on whether the bank relies on credentialed systems, whether the credentials are current, and whether the bank has appropriate governance procedures for monitoring credential status — rather than independently evaluating proprietary AI architectures.

This approach leverages the FDIC's existing supervisory philosophy. The FDIC has long relied on third-party attestations — external audits, SOC reports, fair lending analyses — as inputs to its examination process. Behavioral safety credentialing extends this model to AI systems.

VI. Specific Recommendations

Empty Set LLC recommends the following additions to the GENIUS Act implementing regulations for FDIC-supervised institutions:

1. Require AI Governance Programs for Stablecoin Activities

FDIC-supervised banks engaged in stablecoin activities should be required to establish and maintain AI governance programs consistent with the Treasury FS AI RMF. These programs should address:

- Identification and inventory of all AI systems used in stablecoin operations
- Behavioral safety standards for each system, including error rate tolerances, false-positive rate thresholds, and drift monitoring parameters
- Regular behavioral safety assessments by qualified independent parties
- Documented human override procedures and escalation protocols
- Demographic impact auditing for AI systems making consequential decisions

2. Recognize Third-Party Behavioral Safety Credentialing

The FDIC should recognize behavioral safety credentialing as a permissible mechanism for satisfying AI governance obligations, with criteria for qualifying credentials that include:

- Evaluation against recognized standards (NIST AI RMF, ISO 42001, Treasury FS AI RMF)
- Independence of the credentialing body from the AI vendor
- Time-limited credentials subject to renewal and revocation
- Defined behavioral safety dimensions (accuracy, fairness, drift, robustness, human override)

3. Coordinate AI Governance Standards with State Banking Departments

The FDIC should work with the Conference of State Bank Supervisors (CSBS) to develop consistent AI governance examination criteria for stablecoin activities. Behavioral safety credentialing criteria recognized by the FDIC should be communicated to state banking departments as a recommended standard, reducing the risk of conflicting examination conclusions.

4. Require Demographic Impact Auditing for BSA/AML AI

AI systems used for BSA/AML transaction monitoring in stablecoin operations should undergo regular demographic impact audits assessing whether false-positive rates, alert escalation patterns, and transaction blocking decisions exhibit statistically significant variation by protected characteristics. This is the BSA/AML equivalent of the fair lending analysis the FDIC already requires for credit models.

5. Incorporate AI Governance into CAMELS Assessment

For institutions engaged in stablecoin activities, the quality of the institution's AI governance program — including the behavioral safety credentialing status of its AI systems — should be incorporated into the Management component of the CAMELS rating. This provides a direct supervisory incentive for institutions to invest in AI behavioral safety and creates a consistent evaluation framework across examinations.

6. Develop Supplemental Examination Procedures for AI in Stablecoin Operations

The FDIC should develop and publish supplemental examination procedures specifically addressing AI governance in stablecoin operations. These procedures should address:

- Evaluation of the bank's AI system inventory and governance framework
- Verification of behavioral safety credentials for vendor-provided AI systems
- Assessment of the bank's monitoring of AI behavioral drift and performance
- Review of human override procedures and escalation documentation
- Evaluation of demographic impact audit results
- Assessment of AI-related insurance coverage and operational risk management

7. Coordinate with NIST and Peer Regulators

The FDIC should coordinate with the NIST AI Agent Standards Initiative and with the OCC, Federal Reserve, and NCUA to ensure that behavioral safety credentialing standards for AI systems in stablecoin operations are consistent across the federal banking agencies. Interoperable credentialing

frameworks ensure that vendors serving institutions supervised by different regulators can obtain a single credential recognized by all — reducing costs and complexity for the banking system.

VII. The Systemic Risk Consideration

The FDIC, as the insurer of bank deposits, has a systemic perspective that individual institutions do not. AI behavioral safety failures in stablecoin operations are not just individual institution risks. They are potential systemic risks when AI systems from concentrated vendors are deployed across hundreds or thousands of institutions simultaneously.

A single AI behavioral safety failure — a training data contamination event, an adversarial exploit, a systematic bias that goes undetected — could affect every institution using that vendor's system at the same time. The consequences could include:

- Simultaneous BSA/AML compliance failures across multiple institutions
- Coordinated discriminatory false-positive patterns that trigger public confidence concerns
- Reserve management errors during market stress that cascade across institutions

The FDIC's implementing regulations should require enhanced behavioral safety credentialing for AI systems deployed by vendors that serve a significant share of FDIC-supervised institutions. This concentration-adjusted approach reflects the FDIC's systemic risk mandate and provides early warning of behavioral safety issues before they propagate across the banking system.

VIII. Responses to Specific FDIC Questions

We recognize that this proposed rule focuses on the application and licensing procedures for stablecoin subsidiaries of FDIC-supervised institutions, with substantive prudential standards to follow in a subsequent rulemaking. Our responses to the FDIC's specific questions address what the application evaluation process should require regarding AI governance — ensuring that AI risk is assessed at the gatekeeping stage, not deferred until after a subsidiary is operational.

Question 3 (Additional Information for Evaluating Statutory Factors): The FDIC should require applicants to submit documentation regarding the AI systems their stablecoin subsidiaries intend to deploy for BSA/AML compliance, reserve management, consumer disclosure, and transaction processing functions. Specifically, applicants should provide: (a) an inventory of planned AI systems and their intended functions; (b) evidence that AI vendors maintain governance programs consistent with the Treasury FS AI RMF; (c) behavioral safety credentials from qualified independent bodies for AI systems performing consequential compliance decisions; and (d) the applicant's planned procedures for monitoring AI behavioral drift, responding to AI safety incidents, and escalating AI-generated decisions to human review. This information is directly relevant to evaluating the statutory factors under Section 5(c) — particularly the adequacy of the applicant's risk management framework and the safety and soundness of the proposed operations.

Question 4 (Additional Evaluation Factors Beyond Section 5(c)): The FDIC should consider AI governance quality as an additional evaluation factor. The GENIUS Act's statutory factors focus on financial adequacy, managerial fitness, and operational soundness. AI governance spans all three: an applicant's financial projections are unreliable if the AI systems performing compliance functions are prone to behavioral safety failures that generate enforcement liability;

managerial fitness is incomplete if the applicant's leadership lacks the capacity to oversee AI risk; and operational soundness is undermined if the AI systems driving stablecoin operations have not been independently evaluated for accuracy, fairness, and robustness. Adding AI governance as an explicit evaluation factor — or clarifying that it falls within the existing statutory factors — ensures that the FDIC assesses this risk at the application stage rather than discovering it during the first examination.

Question 5 (Information Demonstrating Capital/Liquidity and Reserve Sufficiency): Beyond the financial information needed to demonstrate reserve adequacy, applicants should be required to submit evidence of AI-related operational risk management. This includes: (a) the applicant's assessment of AI-related insurance coverage — specifically, whether the stablecoin subsidiary's AI-related operational liabilities are covered under commercial policies or will require affirmative AI liability coverage; (b) the applicant's capital allocation for AI-related operational risk, recognizing that the January 2026 ISO exclusions (CG 40 47/48/35 08) may leave AI-related liabilities uninsured; and (c) behavioral safety credentials for AI systems that manage reserves, as reserve management errors during stress events could threaten the subsidiary's ability to maintain the one-to-one reserve ratio required by the GENIUS Act.

Question 6 (Complex Ownership and Consortium Structures): The FDIC correctly anticipates that stablecoin subsidiaries may involve consortium or multi-party structures. When multiple institutions participate in a shared stablecoin infrastructure, the AI systems powering that infrastructure serve all participating institutions simultaneously. The FDIC should require consortium applicants to demonstrate that the shared AI systems have been independently evaluated for behavioral safety and that governance responsibilities for AI risk are clearly allocated among consortium members. A behavioral safety failure in a shared AI system could simultaneously affect multiple FDIC-supervised institutions — a concentration risk that the application process should identify and address.

Question 7 (Policies, Procedures, and Customer Agreements): The FDIC should require that applicants' customer agreements include disclosure of AI system involvement in stablecoin operations. When AI systems process transactions, generate disclosures, or make decisions that affect a customer's access to funds, the customer should know that an AI system is involved, what behavioral safety standards that system meets, and how to request human review of an AI-generated decision. These provisions should be evaluated as part of the application review to ensure that AI transparency is designed into the subsidiary's operations from inception.

Question 11 (Hidden Compliance Costs or Macroeconomic Effects): The FDIC's initial analysis should account for the cost of AI governance compliance. Stablecoin subsidiaries will deploy AI systems that require ongoing behavioral safety evaluation, demographic impact auditing, and continuous monitoring. These costs are real but manageable — particularly when amortized through third-party behavioral safety credentialing that serves multiple institutions through a single vendor evaluation. The alternative — each institution independently validating its vendor's AI systems — is far more costly and produces inconsistent results. The FDIC should recognize behavioral safety credentialing as a cost-reducing compliance mechanism in its regulatory impact analysis.

IX. Cross-Reference: Treasury Section 9 Report

The Treasury's Section 9 report, *"Innovative Technologies to Counter Illicit Finance Involving Digital Assets"* (March 2026), independently validates the concerns raised in this comment. Section 4 of the report identifies AI and machine learning as essential tools for BSA/AML compliance in the stablecoin ecosystem, while acknowledging that AI deployment introduces resource burdens for smaller financial institutions, consumer privacy risks, and cybersecurity vulnerabilities.

Behavioral safety credentialing addresses all three Treasury concerns: it reduces resource burdens by enabling institutions to rely on independent evaluation rather than building internal capacity; it incorporates data provenance auditing as a credential dimension, addressing privacy risk; and it includes adversarial robustness testing, addressing the cybersecurity vulnerability of the AI systems themselves.

The FDIC should align its application evaluation criteria with the Treasury's findings — ensuring that applicants deploying AI systems for stablecoin compliance demonstrate governance practices consistent with the risks the Treasury has identified.

Closing

The FDIC's implementation of the GENIUS Act will determine whether stablecoin activities are conducted safely across the state-chartered banking system. The AI systems that will power stablecoin compliance — transaction monitoring, reserve management, consumer disclosure, regulatory reporting — will exercise consequential authority that directly implicates the FDIC's safety-and-soundness and deposit insurance missions.

Behavioral safety credentialing provides the FDIC with a practical, scalable, and proportionate mechanism for AI governance oversight. It resolves the validation gap between the Treasury FS AI RMF's 230 control objectives and the FDIC's examination capacity. It provides community banks with a path to AI governance that does not require internal model risk management teams. It creates a common standard that can be coordinated with state banking departments through the CSBS. And it generates the governance evidence that the insurance market needs to offer affirmative AI liability coverage — reducing uninsured operational risk on supervised institutions' balance sheets.

Empty Set LLC is developing the credentialing infrastructure that will enable banks, regulators, and insurers to verify the behavioral safety of AI systems in financial services. We urge the FDIC to incorporate behavioral safety credentialing into the GENIUS Act regulations — establishing a framework that protects depositors, supports community banks, and provides examiners with the tools they need to supervise AI risk in the stablecoin era.

We appreciate the opportunity to comment and welcome further dialogue with the FDIC and its staff on these matters.

Respectfully submitted,

Brice Love Co-Founder, Empty Set LLC brice@emptysetllc.com