Public Comment on OCC Notice of Proposed Rulemaking — Stablecoin Activities Under the GENIUS Act

Submitted by: Empty Set LLC (Nevada Series LLC) **Contact:** Brice Love, Co-Founder **Email:** brice@emptysetllc.com **Date:** March 20, 2026 **Re:** Notice of Proposed Rulemaking — Implementing the Guiding and Establishing National Innovation for U.S. Stablecoins Act for the Issuance of Stablecoins by Entities Subject to the Jurisdiction of the Office of the Comptroller of the Currency, 91 FR 10202 (March 2, 2026), Docket No. OCC-2025-0372, RIN 1557-AF41 **Submitted via:** Regulations.gov, https://www.regulations.gov/docket/OCC-2025-0372

Executive Summary

Empty Set LLC respectfully submits this comment in response to the Office of the Comptroller of the Currency's Notice of Proposed Rulemaking implementing the Guiding and Establishing National Innovation for U.S. Stablecoins Act ("GENIUS Act"). We write to address an issue that the proposed rule does not yet adequately confront: the behavioral safety of AI systems that national banks will deploy to comply with the Act's BSA/AML, reserve management, and consumer protection requirements.

This is not a peripheral concern. The operational demands the GENIUS Act places on stablecoin-issuing institutions — continuous reserve attestation, real-time transaction monitoring, sanctions screening at machine speed, and suspicious activity detection across novel asset classes — will require AI-driven automation at a scale that manual compliance programs cannot sustain. National banks entering stablecoin activities will not hire hundreds of additional compliance officers. They will deploy AI agents to monitor transactions, flag suspicious activity, manage reserves, and generate regulatory reports.

The question is not whether AI will be used in GENIUS Act compliance. It is whether the AI systems performing that compliance will themselves be subject to governance standards that ensure they operate safely, accurately, and without discriminatory bias. Today, no such standard exists in the OCC's supervisory framework. The existing model risk management guidance — OCC Bulletin 2011-12 (SR 11-7) and OCC Bulletin 2023-17 on third-party risk management — was designed for traditional statistical models and vendor relationships. Neither contemplates autonomous AI agents that make real-time compliance decisions with minimal human oversight.

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 directly addresses the gap between the Treasury's February 2026 Financial Services AI Risk Management Framework — which establishes 230 control objectives for AI governance in financial services — and the practical examination tools that OCC supervisors need to assess whether a bank's AI systems meet those objectives. We propose that the OCC incorporate behavioral safety credentialing requirements into the GENIUS Act implementing regulations as a mechanism for third-party validation of AI systems used in stablecoin operations.

I. The AI Compliance Imperative in Stablecoin Operations

The GENIUS Act imposes operational requirements on stablecoin issuers that are, by design, continuous and real-time. Unlike traditional banking activities where compliance reviews occur periodically — quarterly call reports, annual BSA/AML examinations, periodic model validation — stablecoin operations demand persistent, automated oversight. Three requirements illustrate why AI deployment is not optional but structural.

A. BSA/AML Transaction Monitoring

The Bank Secrecy Act requires financial institutions to monitor customer transactions for suspicious activity and file Suspicious Activity Reports (SARs) with FinCEN. For traditional banking products, transaction volumes are manageable through a combination of rules-based screening systems and human review of flagged alerts.

Stablecoin transactions operate at fundamentally different velocities. A single stablecoin issuer may process thousands of peer-to-peer transfers per hour, each requiring sanctions screening against OFAC's Specially Designated Nationals list, transaction pattern analysis for layering and structuring, and beneficial ownership verification when transfers cross jurisdictional thresholds. The volume, speed, and pseudonymous nature of blockchain-based transactions make human-primary compliance impractical.

National banks entering stablecoin activities will deploy AI-driven transaction monitoring systems that autonomously evaluate transfers, assign risk scores, escalate alerts, and in some configurations, automatically block or delay transactions pending review. These systems will make consequential compliance decisions — decisions that determine whether a customer's funds move or freeze — with minimal real-time human oversight.

B. Reserve Management and Attestation

The GENIUS Act requires stablecoin issuers to maintain reserves equal to or exceeding the outstanding supply of stablecoins, and to provide periodic attestation of reserve adequacy. For large-scale issuers, reserve management involves continuous rebalancing across eligible asset classes (Treasury securities, insured deposits, repurchase agreements), real-time monitoring of redemption flows, and automated triggers for reserve replenishment.

AI systems are already used in treasury management and asset-liability modeling at major banks. The GENIUS Act's reserve requirements will extend this reliance to stablecoin-specific contexts where the speed of redemption requests — potentially thousands per minute during market stress events — demands algorithmic decision-making about which reserves to liquidate, when to halt redemptions, and how to communicate reserve status to regulators.

C. Consumer Protection and Disclosure

The Act imposes disclosure requirements on stablecoin issuers, including the right of holders to redeem at par, the composition of reserves, and the risks associated with holding stablecoins. Compliance with these requirements at scale — across millions of holders, in real time, through multiple channels — will increasingly be mediated by AI systems that generate disclosures, respond to holder inquiries, and manage redemption processes.

Each of these functions places an AI system in a position of consequential authority: the authority to

freeze a customer's transaction, to liquidate reserves, or to represent the issuer's financial condition to holders. The behavioral safety of these systems — their accuracy, their bias characteristics, their failure modes under stress — is not a secondary consideration. It is the primary determinant of whether GENIUS Act compliance is genuine or performative.

II. The Supervisory Gap: SR 11-7 Was Not Built for This

The OCC's existing model risk management framework, articulated in OCC Bulletin 2011-12 (implementing the interagency SR 11-7 guidance), establishes three pillars of sound MRM practice: model development, model validation, and model governance. This framework has served the banking system well for over a decade. It was designed, however, for a specific class of models: statistical models with defined inputs, quantifiable outputs, and periodic recalibration cycles.

Autonomous AI agents deployed for GENIUS Act compliance differ from traditional models in ways that SR 11-7 does not address:

Continuous adaptation. Traditional models are validated at a point in time and recalibrated on a schedule. AI agents, particularly those using machine learning for transaction monitoring, continuously adapt their behavior based on new data. A model validated in January may behave materially differently by March — not because of a deliberate recalibration, but because the agent has learned from intervening data. SR 11-7's periodic validation framework does not capture this continuous behavioral drift.

Emergent behavior. Traditional models produce outputs that are deterministic functions of their inputs. AI agents, particularly those with generative or reinforcement learning components, can exhibit emergent behaviors — outputs that were not explicitly programmed and may not have been observed during validation. An AI transaction monitoring agent trained on historical SAR data might, under novel market conditions, begin flagging legitimate transactions from specific geographic regions or demographic groups at disproportionate rates. This emergent bias is not a model error in the traditional sense. It is a behavioral safety failure.

Opacity. SR 11-7 requires that models be "conceptually sound" and that their "processing component" be understood. For complex AI systems — deep neural networks, transformer architectures, ensemble methods — the processing component may be mathematically opaque even to the institution deploying it, and certainly to the third-party vendor that developed it. The OCC's third-party risk management guidance (Bulletin 2023-17) requires banks to manage vendor risk, but it does not provide tools for validating the behavioral safety of AI systems whose internal logic is not interpretable by the bank's own risk management staff.

Real-time consequential authority. Traditional models inform human decisions. AI agents in stablecoin compliance will make decisions — blocking transactions, adjusting reserves, generating regulatory filings — that have immediate, material consequences for customers and for the bank's regulatory standing. The speed at which these decisions are executed may preclude meaningful human-in-the-loop review for the majority of transactions.

The proposed GENIUS Act regulations should acknowledge that AI systems deployed for stablecoin compliance require a supervisory framework that supplements SR 11-7 with behavioral safety standards specific to autonomous, continuously adapting, and consequentially empowered AI agents.

III. The Discriminatory False-Positive Problem

The most acute behavioral safety risk in AI-driven BSA/AML compliance is not that the system will miss illicit activity — though that risk is real. It is that the system will systematically misidentify legitimate activity as suspicious, and that these false positives will disproportionately burden specific demographic groups, geographic regions, or business categories.

This is not a theoretical concern. The existing BSA/AML compliance ecosystem already produces documented disparate impact. Studies of traditional transaction monitoring systems have found that communities with higher concentrations of cash-intensive small businesses — which correlate with minority and immigrant communities — generate disproportionate SAR filings. The introduction of AI-driven monitoring systems trained on this historically biased data does not eliminate the bias. It automates and scales it.

In the stablecoin context, the consequences of discriminatory false positives are more severe than in traditional banking because of the speed and finality of automated enforcement. When an AI agent flags a stablecoin transfer as suspicious and automatically delays or blocks the transaction, the affected customer experiences an immediate and potentially cascading disruption: failed payments, missed settlement deadlines, reputational damage in decentralized networks where transaction histories are publicly visible.

The OCC's implementing regulations should require that AI systems used for GENIUS Act BSA/AML compliance undergo regular demographic impact audits — testing whether the system's false-positive rates vary by race, ethnicity, geography, business type, or other protected characteristics. This is the behavioral safety equivalent of the fair lending analysis that the OCC already requires for credit models. The difference is that BSA/AML models have historically escaped this scrutiny because they are framed as law enforcement tools rather than consumer-facing decisions. The GENIUS Act's consumer protection provisions create an opportunity to close this gap.

IV. 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), represents the most comprehensive attempt to date to operationalize AI governance standards for the financial sector. Its 230 control objectives span governance, data management, model development, validation, monitoring, third-party risk, and consumer protection. The framework adapts the NIST AI Risk Management Framework to the specific regulatory, operational, and fiduciary context of financial services.

The FS AI RMF is the right standard. The problem is the validation mechanism.

The framework establishes what banks should do. It does not establish how regulators — or banks themselves — verify that AI systems actually comply. OCC examiners reviewing a national bank's stablecoin AI systems will need to assess compliance with 230 control objectives. Many of these objectives involve technical evaluations — model performance metrics, drift analysis, data provenance audits, adversarial robustness testing — that exceed the current skill sets of even specialized examination teams.

This is not a criticism of OCC staff. It is a structural observation. The complexity and opacity of

modern AI systems create an information asymmetry between supervised institutions (who deploy the systems) and supervisors (who must evaluate them). This asymmetry is compounded by the Confidential Supervisory Information (CSI) paradox: banks cannot share examination findings with their AI vendors, and vendors cannot disclose proprietary model architectures to regulators without risking competitive harm. The result is that neither the bank nor the examiner may have a complete picture of the AI system's behavioral characteristics.

Third-party behavioral safety credentialing resolves this information asymmetry. An independent credentialing body evaluates the AI system against defined behavioral safety standards — error rate tolerance, discriminatory false-positive rates, drift characteristics, human override latency, adversarial robustness — and issues a verifiable credential that can be shared with the bank, the regulator, and the bank's insurer without disclosing proprietary model architecture.

This is the functional equivalent of a SOC 2 audit for AI systems. The credentialing body attests that the system meets defined behavioral standards. The bank relies on the attestation for its model risk management obligations. The examiner relies on the attestation for its supervisory assessment. The insurer relies on the attestation for its underwriting decisions. No party needs access to the proprietary model architecture. Each party gets the information it needs in a form it can act on.

V. The Insurance Dimension: Why Underwriters Care

The OCC may reasonably ask why a standards development initiative focused on insurance credentialing is commenting on stablecoin regulation. The answer is that the insurance market's evaluation of AI risk has direct implications for the safety and soundness of banks entering stablecoin activities.

On January 1, 2026, the Insurance Services Office released endorsements CG 40 47 (Artificial Intelligence Exclusion), CG 40 48 (Artificial Intelligence Limited Coverage), and CG 35 08 (Artificial Intelligence Amendatory Endorsement) for commercial general liability policies. These endorsements — now being adopted across the property-casualty market — exclude or restrict coverage for liabilities arising from generative AI systems. The practical effect is that AI-related operational liabilities are increasingly uninsurable under standard commercial policies.

For national banks deploying AI in stablecoin operations, this creates a safety-and-soundness concern. Operational risk self-insurance reserves must now account for AI-related liabilities that were previously covered by commercial policies. Banks that cannot demonstrate behavioral safety credentialing for their AI systems will face higher insurance costs, reduced coverage, and potentially uninsurable operational risk exposures — all of which affect capital adequacy and risk management assessments.

The OCC's implementing regulations should recognize that behavioral safety credentialing of AI systems is not merely a best practice for vendor management. It is a precondition for maintaining adequate risk transfer capacity in the emerging AI liability insurance market. Banks that invest in credentialing their AI systems will have access to affirmative AI liability coverage. Banks that do not will bear the full cost of AI-related operational failures on their balance sheets.

VI. Specific Recommendations

Empty Set LLC recommends the following additions to the GENIUS Act implementing regulations:

1. Require AI Governance Programs for Stablecoin Activities

The proposed rule should require national banks engaged in stablecoin activities to establish and maintain AI governance programs that address the behavioral safety of all AI systems used in BSA/AML compliance, reserve management, consumer disclosure, and transaction processing. These programs should be consistent with the Treasury FS AI RMF and should include, at a minimum:

- 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 (transaction blocking, account restriction, SAR generation)

2. Recognize Third-Party Behavioral Safety Credentialing

The proposed rule should recognize third-party behavioral safety credentialing as a permissible mechanism for satisfying AI governance obligations, analogous to the role of external audits in financial reporting or SOC 2 attestations in information security. Specifically:

- Banks should be permitted to rely on credentials issued by qualified independent bodies that evaluate AI systems against recognized standards (NIST AI RMF, ISO 42001, Treasury FS AI RMF)
- Credentials should be verifiable, time-limited, and subject to revocation upon material behavioral safety failures
- The OCC should maintain a list of recognized credentialing frameworks — or delegate this function to NIST — to ensure consistency and prevent credential proliferation

3. Address Continuous Monitoring and Behavioral Drift

The proposed rule should require continuous monitoring of AI systems used in stablecoin operations, not merely periodic validation. Specifically:

- AI systems should be subject to ongoing behavioral drift analysis, with defined thresholds that trigger human review, model revalidation, or system deactivation
- Monitoring data should be available to OCC examiners upon request
- Material changes in AI system behavior — including changes resulting from continuous learning rather than deliberate recalibration — should be reported as model changes under existing MRM guidance

4. Require Demographic Impact Auditing for BSA/AML AI

The proposed rule should require that AI systems used for BSA/AML transaction monitoring in stablecoin operations undergo regular demographic impact audits. These audits should assess whether the system's false-positive rates, alert escalation patterns, and transaction blocking decisions exhibit statistically significant variation by race, ethnicity, geography, or business category. Results should be documented and available to OCC examiners.

5. Integrate AI Risk into Stablecoin Examination Procedures

The OCC should develop supplemental examination procedures for evaluating the behavioral safety of AI systems in stablecoin operations. These procedures should address:

- The adequacy of the bank's AI governance program
- The quality and independence of behavioral safety assessments
- The bank's monitoring of AI behavioral drift
- The bank's procedures for human override of AI compliance decisions
- The bank's assessment of AI-related insurance coverage and operational risk capital

6. Coordinate with NIST AI Agent Standards Initiative

The NIST AI Agent Standards Initiative, launched in February 2026, is developing frameworks for AI agent identity verification, authorization, and interoperability. The OCC should coordinate with NIST to ensure that AI systems used in stablecoin operations can be uniquely identified, authenticated, and monitored across institutions — enabling systemic risk surveillance and facilitating information sharing about behavioral safety failures without disclosing proprietary model architectures.

VII. The Cross-Sector Convergence

The AI governance gap in stablecoin regulation is not unique to the OCC. It is the same gap that the FDIC faces in its GENIUS Act implementation, that the NCUA faces for federally insured credit unions, that the FHFA confronted in Freddie Mac's Section 1302.8 AI governance requirements, and that state insurance commissioners are grappling with as they respond to the ISO CG 40 47 exclusion wave.

Every financial regulator is converging on the same structural problem: institutions must validate AI systems they procured from vendors but cannot access the underlying model architecture. The comments submitted in response to the OSTP's Request for Information on AI in financial services confirmed this as the industry's primary concern. The Confidential Supervisory Information rules prevent examination findings from being shared with vendors. Proprietary model architectures prevent vendors from providing full transparency to regulators. The result is an information vacuum where neither the institution, the regulator, nor the insurer has a complete picture of the AI system's behavioral characteristics.

Behavioral safety credentialing fills this vacuum. It provides a standardized, verifiable, and shareable attestation of AI system performance that serves all three parties without requiring disclosure of proprietary information. The OCC has an opportunity to establish this framework in the GENIUS Act regulations — creating a precedent that other regulators can adopt and that the industry can build upon.

VIII. Responses to Specific OCC Questions

This comment is responsive to several of the 211 numbered questions the OCC has posed for public input. We address the most directly relevant here:

Question 107 (Consumer Protection and Disclosures): The proposed disclosure categories are necessary but insufficient. They address financial disclosures — reserve composition, redemption rights, fee schedules — but do not address the behavioral characteristics of the AI systems that will mediate consumers' interactions with stablecoin services. When an AI agent processes a consumer's redemption request, generates a disclosure, or blocks a transaction pending review, the consumer's experience is determined by the AI system's behavioral safety profile, not by the disclosure text alone. The OCC should require that issuers disclose whether AI systems used in consumer-facing stablecoin operations have been evaluated against recognized behavioral safety standards by qualified independent parties, and the results of such evaluations. This is the AI equivalent of the cybersecurity disclosure that consumers now expect from financial institutions.

Question 198 (Systemic Risk and Financial Stability): The proposed regulatory architecture does not yet address the systemic risk created by AI systems deployed across multiple stablecoin issuers. When multiple national banks procure AI-driven compliance systems from the same concentrated set of vendors, a behavioral safety failure in a single vendor's AI system — a training data contamination, an adversarial exploit, a systematic bias — could simultaneously compromise BSA/AML compliance across the stablecoin ecosystem. The OCC should mandate the collection of AI system telemetry data — behavioral drift indices, false-positive rates, override frequencies — as part of the additional economic data referenced in this question. Furthermore, the OCC should require that AI systems used in stablecoin operations carry verifiable behavioral safety credentials, enabling systemic monitoring of AI governance quality across the supervised population.

Question 201 (Size-Based Regulatory Adjustments): Behavioral safety standards for AI systems should not be adjusted based on issuer size. A discriminatory false-positive that freezes a customer's stablecoin transaction is equally harmful regardless of whether the issuer has $100 million or $100 billion in outstanding issuance. However, the *mechanism* for demonstrating AI behavioral safety should be proportionate. Large issuers with internal model risk management teams may satisfy AI governance obligations through internal assessments. Smaller issuers should be permitted to satisfy the same obligations through third-party behavioral safety credentialing — which places the evaluation burden on an independent body rather than on the issuer's internal staff. Size should determine the *how*, not the *whether*, of AI governance compliance.

Question 203 (Foreign Jurisdictions and On-Chain Identification): The question of how stablecoin holders can technically distinguish GENIUS Act-compliant tokens from non-compliant tokens on a public blockchain has a natural extension: how can regulators, insurers, and consumers distinguish between stablecoin issuers whose AI compliance systems meet behavioral safety standards and those that do not? Behavioral safety credentials — verifiable, time-limited, and revocable attestations of AI system performance — can be anchored to on-chain metadata or linked to the issuer's regulatory filings, providing a transparent, machine-readable signal of AI governance quality alongside the compliance status of the stablecoin itself.

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 (pages 12-16) identifies AI and machine learning as essential tools for BSA/AML compliance in the stablecoin ecosystem, while simultaneously acknowledging that AI deployment "introduces novel resource burdens for smaller financial institutions, raises significant consumer

privacy risks regarding the data ingested by the models, and creates new cybersecurity vulnerabilities if the AI systems themselves are compromised by malicious actors."

Behavioral safety credentialing directly addresses all three of the Treasury's concerns: it reduces the resource burden by allowing institutions to rely on independent credentialing rather than building internal evaluation 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.

Closing

The GENIUS Act represents a significant expansion of the national banking system's engagement with digital assets. The operational infrastructure required to support stablecoin activities — continuous transaction monitoring, real-time reserve management, automated consumer disclosures — will be built on AI systems that exercise consequential authority over customer funds, regulatory compliance, and institutional risk.

The OCC's implementing regulations should ensure that the behavioral safety of these AI systems receives the same supervisory attention as the financial safety of the reserves backing the stablecoins. The tools for this oversight exist: the Treasury FS AI RMF provides the standard, third-party behavioral safety credentialing provides the validation mechanism, and NIST's AI Agent Standards Initiative provides the interoperability framework.

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 OCC to incorporate behavioral safety credentialing into the GENIUS Act regulations — not as a prescriptive mandate, but as a recognized mechanism for demonstrating AI governance compliance that benefits institutions, supervisors, and consumers alike.

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

Respectfully submitted,

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