Public Comment on SEC Petition for Rulemaking 4-882

Submitted by: Empty Set LLC (Nevada Series LLC) **Contact:** Brice Love, Co-Founder **Email:** brice@emptysetllc.com **Date:** March 10, 2026 **Re:** Petition for Rulemaking 4-882 — Mandatory AI Governance and Risk Management Disclosure (filed by Candace Arthur, February 9, 2026)

Executive Summary

Empty Set LLC respectfully submits this comment in support of Petition for Rulemaking 4-882, filed by Candace M. Arthur on February 9, 2026, which proposes mandatory AI governance and risk management disclosures modeled on the Commission's 2023 cybersecurity disclosure framework. Ms. Arthur's petition — grounded in her restructuring practice at Latham & Watkins — correctly frames unguarded AI deployment as a capital preservation problem, not merely an ethics concern. Algorithmic failures trigger debt covenants, subordinate equity, and destroy enterprise value. The petition's proposed Generally Accepted Algorithmic Principles (GAIP) framework, with its ten foundational principles mapped from GAAP and three implementation pillars, provides a structured vocabulary for this emerging category of material risk.

We write to strengthen the GAIP framework with a specific, actionable recommendation: disclosure requirements should encompass not only governance *processes* but also behavioral safety *metrics* — quantifiable indicators of whether an issuer's AI systems operate within defined safety boundaries under real-world conditions. The petition's concept of "integrity as a mandatory cost variable within the AI's optimization function" captures the right principle. Behavioral safety credentialing provides the measurement infrastructure to make that principle auditable and comparable across issuers.

Empty Set LLC operates a standards development initiative focused on behavioral safety credentialing for autonomous AI agents, with particular emphasis on the insurance and financial services sectors. Our work sits at the intersection of the three forces converging on public companies today: (1) the insurance industry's withdrawal of coverage for AI-related liabilities, (2) courts' extension of strict products liability to AI agent behavior, and (3) state legislatures' imposition of mandatory behavioral safety requirements. Each of these forces generates material risk that current SEC disclosure rules do not adequately capture.

The Commission has precedent for exactly this kind of intervention. In July 2023, the SEC adopted final rules requiring registrants to disclose material cybersecurity incidents and to describe their cybersecurity risk management, strategy, and governance on an annual basis. That rulemaking recognized that cybersecurity risk had matured from a technical concern into a material investor concern — and that voluntary disclosure was producing inconsistent, incomparable information. AI behavioral safety risk has reached the same inflection point. The Investor Advisory Committee's December 4, 2025 recommendation — urging the Commission to adopt AI disclosure guidelines covering AI definitions, board oversight, and separate impact reporting — confirms that this inflection point is recognized within the Commission's own advisory structure. The insurance market data, judicial rulings, and legislative activity we describe below demonstrate that this is no longer a speculative concern. It is a present, quantifiable, and rapidly compounding source of financial exposure for public companies.

We recognize that Chair Atkins has articulated a principles-based, technology-neutral regulatory philosophy that favors materiality-driven disclosure over prescriptive mandates. We believe the

recommendations in this comment are fully compatible with that philosophy. We do not propose new categories of regulation. We propose that the Commission incorporate AI governance into the existing cybersecurity disclosure framework through interpretive guidance — leveraging the structural template the Commission has already built — and that behavioral safety credentialing be recognized as one indicator of governance maturity within that framework, just as SOC 2 attestations and ISO 27001 certifications function in cybersecurity disclosure today.

I. Why Behavioral Safety Metrics Matter for Investor Disclosure

The petition's GAIP framework proposes disclosure obligations that mirror the cybersecurity rule's structure: risk management processes (analogous to Reg S-K Item 106(b)), governance oversight (analogous to Item 106(c)), material incident reporting (analogous to Form 8-K Item 1.05), and integration into business description and legal proceedings narratives (analogous to Items 101 and 103). This structural parallel is well-conceived and reduces the implementation burden for both issuers and Commission staff. The ten GAIP principles — regularity, consistency, sincerity, permanence of methods, non-compensation, prudence, continuity, periodicity, materiality, and utmost good faith — provide a principled vocabulary for evaluating AI governance quality.

However, governance process disclosure alone is insufficient. An issuer can maintain an AI ethics board, publish responsible AI principles, and designate a Chief AI Officer while deploying systems whose behavioral characteristics create substantial, undisclosed financial exposure. The GAIP principle of "prudence" — validation and testing before deployment — and the principle of "continuity" — ongoing monitoring and reassessment — both imply the existence of measurable outcomes. Behavioral safety credentialing makes those outcomes explicit and comparable.

Behavioral safety credentialing provides investors with verifiable, comparable data on whether an AI system's outputs remain within defined operational boundaries. This is not an abstract concept. It is the functional equivalent of what the insurance industry calls "loss control" — the measurable practices that determine whether a risk is insurable and at what premium. When an insurer evaluates whether to underwrite a commercial property, it does not ask whether the building owner has a fire safety committee. It asks whether the building has sprinklers, what their flow rate is, and when they were last inspected. AI behavioral safety credentialing applies the same logic to autonomous systems: not whether governance exists, but whether it produces measurable safety outcomes.

The distinction matters for investors because governance process disclosure creates a compliance floor without a performance signal. Every Fortune 500 company will eventually have an AI governance framework. The differentiating question — the question that drives investment risk — is whether that framework produces AI systems that behave predictably under adversarial, edge-case, and high-stakes conditions. Behavioral safety metrics answer that question. Governance committee rosters do not.

The SEC's 2023 cybersecurity disclosure rule implicitly recognized this principle. While the rule requires disclosure of governance processes (board oversight, management's role), it also requires disclosure of whether the registrant engages assessors, consultants, auditors, or other third parties in connection with cybersecurity risk management. This provision created a natural disclosure incentive for companies to obtain third-party cybersecurity certifications — SOC 2, ISO 27001, NIST CSF — because those certifications provide concrete, affirmative content for the required

disclosure. AI behavioral safety credentialing should occupy the same structural role in any AI governance disclosure framework.

II. The Insurance Market Is Already Pricing This Risk

The most powerful evidence that AI behavioral safety is a material investor concern comes not from regulators or academics but from the insurance industry itself — the market whose entire business model depends on accurate risk quantification.

The scale of exposure is already quantified. The National Association of Insurance Commissioners' 2025 survey across sixteen states found that 84% of health insurers, 88% of auto insurers, and 70% of homeowners insurers are deploying AI or machine learning systems in material business operations. Yet fewer than half of those insurers have implemented formal bias testing, vendor management, or explainability documentation for their AI systems. The governance gap is not hypothetical — the industry's own data proves it.

On January 1, 2026, the Insurance Services Office (ISO), a Verisk Analytics subsidiary that develops the standard policy forms used by approximately 70% of the U.S. property-casualty market, released three new endorsements that fundamentally restructured AI liability coverage in commercial general liability (CGL) policies:

- **CG 40 47** — Artificial Intelligence Exclusion: Excludes bodily injury and property damage arising out of AI systems from standard CGL coverage.
- **CG 40 48** — Artificial Intelligence Limited Coverage: Provides a narrow, sublimited exception to the exclusion for specified AI use cases, subject to additional underwriting.
- **CG 35 08** — Artificial Intelligence Amendatory Endorsement: Modifies existing policy definitions to address AI-specific liability scenarios.

The practical effect is that the default coverage position for AI-related liabilities under standard CGL policies has shifted from *included* to *excluded*. Companies deploying AI systems must now affirmatively seek coverage, submit to additional underwriting scrutiny, and in many cases accept sublimits, higher retentions, or outright declinations.

This is not a gradual tightening. It is a structural withdrawal of capacity. The managing general agent (MGA) market — the specialty underwriters that had been writing AI liability coverage on a surplus lines basis — is simultaneously contracting as reinsurers demand behavioral safety data that most insureds cannot provide. The result is a coverage gap that represents a direct, quantifiable, and largely undisclosed financial exposure for any public company deploying AI at scale.

Investors reviewing current SEC filings have almost no visibility into this exposure. A registrant's risk factor disclosures may reference "risks associated with artificial intelligence" in general terms, but they rarely disclose whether the company's AI-related liabilities are actually covered by insurance, whether coverage has been reduced or excluded, or what behavioral safety standards the company meets — or fails to meet — in the eyes of its underwriters.

The Commission should recognize that the insurance industry's pricing decisions constitute a market-validated materiality signal. When the industry that professionally quantifies risk concludes that AI behavioral safety is uninsurable without specific credentialing data, that conclusion carries evidentiary weight for securities disclosure purposes.

III. The Legal Landscape Is Accelerating

The judiciary and state legislatures are independently confirming that AI behavioral safety creates legally cognizable liability — liability that registrants must disclose to investors.

Judicial Developments

Garcia v. Character Technologies established that operators of conversational AI systems can face liability under a products liability framework when the system's behavioral outputs cause foreseeable harm. The court's analysis focused not on whether the defendant had AI governance policies, but on whether the system's *behavior* deviated from what a reasonable user would expect — a behavioral safety standard.

Gavalas v. Google extended this reasoning to AI systems integrated into consumer-facing products, holding that the deployer of an AI agent bears responsibility for the agent's behavioral characteristics in the same manner that a product manufacturer bears responsibility for design defects. The court explicitly rejected the argument that disclosure of AI's "experimental" nature immunized the deployer from liability for behavioral failures.

These cases signal a judicial trajectory toward strict liability for AI behavioral harms — meaning that governance intentions and disclosure disclaimers will not shield issuers from liability. What will matter is whether the AI system's behavior met an objective safety standard. This is precisely the information that behavioral safety credentialing generates and that investors need to assess litigation risk.

State Legislative Activity

Multiple states have enacted or are advancing legislation that imposes specific behavioral safety requirements on AI deployers:

- **California SB 243** mandates behavioral risk assessments for high-risk AI systems and requires deployers to implement ongoing monitoring of AI system outputs against defined safety benchmarks.
- **Colorado SB 24-205** (the Colorado AI Act) requires deployers of high-risk AI systems to implement risk management policies, conduct impact assessments, and disclose the system's behavioral characteristics to affected individuals.
- **New York RAISE Act** proposes mandatory behavioral auditing for AI systems used in consequential decisions, with public reporting requirements.

Each of these statutes creates compliance obligations and potential penalties that constitute material risk for registrants operating in those jurisdictions. More importantly, they establish a patchwork of state-level behavioral safety standards that public companies must navigate — precisely the kind of fragmented regulatory landscape that federal disclosure requirements can help investors understand and compare across issuers.

IV. Specific Recommendations

The most productive path forward may not be formal rulemaking — which faces legitimate concerns about prescriptive mandates in a rapidly evolving technology space — but interpretive guidance that incorporates AI governance into the Commission's existing cybersecurity disclosure framework. The structural template already exists. The following recommendations are designed to be implementable through staff guidance, comment letters, and examination priorities without requiring new rules or categories of regulation.

Empty Set LLC recommends the following:

1. **Interpretive guidance clarifying that AI governance falls within the existing cybersecurity disclosure framework.** The 2023 cybersecurity rule's requirements — risk management process disclosure (Reg S-K Item 106(b)), governance oversight disclosure (Item 106(c)), and material incident reporting (Form 8-K Item 1.05) — are structurally sufficient to encompass AI-related risks. Staff guidance clarifying that "cybersecurity" encompasses algorithmic and AI system risks where those systems are material to registrant operations would extend the framework without new rulemaking. The petition's proposed GAIP disclosure structure maps directly to this existing architecture.

2. **Encourage disclosure of behavioral safety credentialing status within existing risk management narratives.** Registrants deploying AI systems in material business operations should be encouraged to disclose whether those systems have been evaluated against recognized behavioral safety standards — such as the NIST AI Risk Management Framework, ISO 42001, or the Treasury Financial Services AI Risk Management Framework (February 2026) — by qualified third parties, and the results of such evaluations. This parallels the cybersecurity rule's existing provision for disclosure of third-party assessor engagement and creates a natural market incentive for issuers to obtain credentialing. The approach is framework-neutral, avoids regulatory lock-in, and allows the market to develop and refine credentialing standards — consistent with Chair Atkins' principles-based regulatory philosophy.

3. **Clarify that AI-related insurance coverage changes are material events warranting disclosure.** The application of Absolute AI Exclusions to a registrant's commercial general liability policies — including endorsements such as ISO CG 40 47/48 and CG 35 08 — constitutes a material reduction in risk transfer capacity. When risk cannot be transferred via insurance, it is borne entirely by the balance sheet and ultimately by shareholders. The Commission's existing materiality framework already captures this, but staff guidance or a sample comment letter would signal that examiners expect this disclosure and that boilerplate AI risk factors are insufficient.

4. **Encourage disaggregated AI risk reporting.** Consistent with the petition's proposal and the IAC's December 2025 recommendation for separate AI impact reporting, registrants should distinguish between AI deployed in internal operations (where risks accrue primarily through operational disruption and workforce effects) and AI deployed in consumer-facing products (where risks accrue through products liability, regulatory action, and reputational harm). This disaggregation improves comparability without imposing prescriptive categories — issuers retain discretion over how to characterize their AI deployments while providing investors with the granularity needed to assess risk.

5. **Establish examination priorities that incentivize good-faith behavioral safety disclosure.** The Division of Examinations' 2026 priorities already include AI accuracy review.

Extending this focus to evaluate the quality and specificity of AI governance disclosures — rewarding registrants that provide concrete, verifiable information over boilerplate — would create a de facto disclosure standard through accumulated staff positions. This incremental, enforcement-driven approach aligns with the Commission's historical practice of using examination priorities and comment letters to shape disclosure norms before (or instead of) formal rulemaking.

Closing

Ms. Arthur's petition arrives at a moment of convergence. The insurance industry is withdrawing coverage. Courts are extending products liability to AI behavior. State legislatures are imposing behavioral safety mandates. The IAC has recommended AI disclosure guidelines. The Treasury has published a financial services AI risk management framework with 230 control objectives. And the Commission's own examination priorities have flagged AI accuracy as a focus area.

The petition's GAIP framework correctly identifies the disclosure gap. The Commission's existing cybersecurity rule provides the structural template to close it. The recommendations in this comment are designed to bridge the two — extending the cybersecurity framework to encompass AI governance risk through interpretive guidance rather than prescriptive mandates, and recognizing behavioral safety credentialing as a market-developed governance indicator rather than a regulatory requirement.

Empty Set LLC is developing the credentialing infrastructure that will allow the insurance market to underwrite AI behavioral risk and that will give investors the comparable, verifiable data they need to assess AI governance quality — not just AI governance existence. We urge the Commission to act on Petition 4-882, and to do so in a manner consistent with its principles-based approach: clarifying that AI governance falls within the existing disclosure framework, encouraging the use of recognized standards, and allowing market-driven credentialing to develop the measurement infrastructure that investors need.

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

Respectfully submitted,

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