SUBJECT: Public Comment on Proposed GSAR 552.239-7001 — "Basic Safeguarding of Artificial Intelligence Systems"

Submitted by: Brice Love, Managing Member, Empty Set LLC | emptysetllc.com | brice@emptysetllc.com

Submitted to: General Services Administration, Federal Acquisition Service — MAS Program Management Office

Date: March 25, 2026

Submitted via: buy.gsa.gov/interact

I. Organizational Identity and Interest

Empty Set LLC is a technology company developing legal, insurance, and credentialing infrastructure for autonomous AI systems. We hold a portfolio of patents covering AI behavioral safety frameworks, credentialing architectures, and autonomous agent governance systems. We are actively engaged with NIST CAISI on AI agent identity and authorization standards, and we submitted comments to the NIST AI 800-2 benchmark evaluation process (acknowledged March 23, 2026 by CAISI).

We write to offer constructive input on proposed GSAR 552.239-7001 from the perspective of a company building the standards and credentialing infrastructure that would make this clause's compliance objectives achievable in practice.

II. General Assessment

We commend GSA for taking a first-mover position on AI governance in federal procurement. The clause addresses real and urgent risks — data sovereignty, supply chain accountability, incident response, and the need for human oversight of agentic AI systems. These are the right problems to solve.

However, the clause as drafted creates significant implementation gaps that will undermine its own objectives. The core problem is structural: the clause imposes comprehensive compliance obligations on contractors and their service providers without establishing any mechanism for standardized, verifiable compliance demonstration. The result is a regime where compliance is asserted but not credentialed — leaving both contractors and the government without a reliable framework for distinguishing compliant AI systems from non-compliant ones.

We offer six specific recommendations below.

III. Recommendation 1: Establish a Third-Party Credentialing Pathway for AI System Compliance

Clause Sections Affected: (e)(6) Documentation, (i) Evaluation

The clause requires contractors to provide NIST AI RMF-consistent documentation on request and grants the government authority to independently evaluate AI systems using its own benchmarks — benchmarks the government has no obligation to disclose. This creates a compliance environment where contractors cannot proactively demonstrate compliance because the evaluation criteria are

opaque, and the government cannot efficiently evaluate across hundreds of MAS contractors using ad hoc benchmarks for each.

We recommend that the clause explicitly recognize third-party AI credentialing bodies — analogous to the Third-Party Assessment Organizations (3PAOs) in the FedRAMP framework — as an accepted mechanism for demonstrating clause compliance. A contractor whose AI systems hold a recognized behavioral safety credential from an accredited standards body should be presumed compliant with sections (e)(6) and (i), subject to government audit rights. This would:

- Give contractors a clear, actionable compliance pathway
- Give the government a scalable evaluation mechanism rather than requiring custom assessments for every MAS contractor
- Align with NIST AI RMF principles, which contemplate third-party conformity assessment
- Reduce the risk of inconsistent enforcement across contracting officers

The FedRAMP model demonstrates that this approach works for complex technical compliance: independent assessors evaluate against published standards, and the government retains audit authority. AI system compliance is structurally analogous.

IV. Recommendation 2: Define "American AI Systems" Using Auditable, Standards-Based Criteria

Clause Sections Affected: (a) Definitions, (e)(2)

Section (e)(2) prohibits the use of "foreign AI systems or any AI components manufactured, developed, or controlled by non-U.S. entities" but does not define what constitutes domestic development in the context of modern AI systems. Contemporary AI systems are built using globally sourced training data, open-source components maintained by international contributors, cloud infrastructure spanning multiple jurisdictions, and multinational development teams. A geographic origin test — "developed and produced in the United States" — is functionally indeterminate for software-based AI systems.

We recommend that "American AI Systems" be defined through auditable compliance criteria rather than geographic origin assertions. Specifically:

- Align with Trade Agreements Act (TAA) standards already used in GSA Schedule procurement for determining country of origin
- Establish verifiable criteria such as: corporate domicile, location of training infrastructure, jurisdiction governing data handling, and whether the system has been modified to comply with any non-U.S. government's content or surveillance requirements
- Allow compliance to be demonstrated through third-party credentialing rather than contractor self-certification

This approach preserves the clause's policy objective — ensuring AI systems serving the federal government are not subject to foreign government control — while providing workable compliance criteria for contractors.

V. Recommendation 3: Allow Credentialed Compliance to Satisfy Service Provider Flow-Down Obligations

Clause Sections Affected: (b) Supremacy, (c) Service Provider Compliance

The clause makes the prime contractor responsible for all service provider compliance, but prime contractors — particularly small and mid-size integrators — have no contractual leverage over upstream AI platform providers who are not parties to the government contract. A small integrator building on a foundation model from a major AI provider cannot compel that provider to comply with the clause's data handling, incident reporting, or American AI requirements. Yet the prime contractor bears full liability, including potential False Claims Act exposure, for the service provider's noncompliance.

We recommend that the clause establish a credentialing-based safe harbor for service provider compliance: if an upstream service provider holds a recognized AI behavioral safety credential from an accredited standards body, the prime contractor's flow-down obligation with respect to that provider should be presumed satisfied. This would:

- Solve the "David and Goliath" problem where small integrators are liable for the compliance of providers orders of magnitude larger
- Create market incentives for major AI platform providers to obtain credentialing
- Provide the government with greater actual assurance than the current system of contractual flow-down to non-parties

VI. Recommendation 4: Distinguish Behavioral Safety Guardrails from Ideological Content Moderation

Clause Sections Affected: (d)(2)(ii), (i)(1)

Section (d)(2)(ii) states that AI systems "must not refuse to produce data outputs or conduct analyses based on the Contractor's or Service Provider's discretionary policies." Section (i)(1) simultaneously requires the system to be "truthful" and to implement "continuous improvement processes to enhance detection and mitigation of performance, trustworthiness, bias."

These requirements are in direct tension. Many output refusals are not ideological content moderation — they are behavioral safety guardrails that protect accuracy, prevent hallucination, flag uncertainty, and avoid generating outputs that the system cannot produce reliably. Removing all refusal capability in the name of eliminating ideological bias would simultaneously eliminate the safety mechanisms that make AI systems trustworthy for government use.

We recommend that the clause draw an explicit distinction between:

- **Ideological content moderation** (which the clause appropriately seeks to prevent in government contexts), and
- **Behavioral safety guardrails** (which protect the government by preventing unreliable, hallucinated, or harmful outputs)

This distinction could be operationalized through reference to credentialed behavioral safety standards: an AI system's safety-motivated refusals are appropriate if they are consistent with published behavioral safety criteria from an accredited standards body. This gives the government a principled and auditable line between prohibited ideological filtering and required safety behavior.

VII. Recommendation 5: Strengthen Agentic AI Governance Provisions

Clause Section Affected: (e)(3) Human Oversight

The clause acknowledges AI systems that use "reasoning, retrieval, or agentic processes" and requires them to "summarize intermediary steps, model routing decisions, and data retrieval methods with full source attribution." This is a meaningful transparency requirement, but it is insufficient for the governance challenges posed by autonomous AI agents operating in government environments.

Agentic AI systems — those that plan, execute multi-step tasks, use tools, and take actions with real-world consequences — require governance mechanisms beyond transparency logging. **We recommend** that the clause establish:

- **Behavioral envelope requirements**: Agentic AI systems should operate within defined behavioral boundaries that specify what actions the system is authorized to take, what decisions require human escalation, and what conditions trigger automatic suspension
- **Decision authority boundaries**: Clear delineation of which decisions the AI system can make autonomously versus which require human approval
- **Auditable action logs**: Not just summarization of reasoning, but complete, tamper-evident logs of all actions taken, including tool use, external API calls, and data access

These requirements would align with emerging best practices in AI agent governance and with the NIST AI RMF's treatment of autonomous systems.

VIII. Recommendation 6: Pursue Formal Rulemaking for a Clause of This Magnitude

The proposed clause would apply to every AI system procured through the GSA MAS program — the government's primary commercial procurement vehicle. Yet it was issued through the MAS Refresh process rather than through formal Federal Acquisition Regulation (FAR) rulemaking, resulting in a compressed comment window and no Federal Register publication.

We recommend that GSA pursue formal rulemaking under the FAR process for a clause of this scope and impact. The compressed timeline has generated significant industry concern and limits the quality of public input. Formal rulemaking would provide broader stakeholder engagement, interagency coordination, and a more durable legal foundation for the clause's requirements. The current MAS Refresh mechanism is appropriate for routine contract updates — not for establishing the government's foundational AI governance framework.

IX. Conclusion

Empty Set LLC supports the policy objectives of GSAR 552.239-7001 and appreciates the opportunity to comment. The clause addresses real governance gaps in federal AI procurement. Our recommendations are intended to make the clause's objectives achievable in practice by establishing credentialing-based compliance pathways, clarifying ambiguous definitions, and strengthening agentic AI governance provisions.

We welcome the opportunity to engage further with GSA on these issues.

Respectfully submitted,

Brice Love Managing Member, Empty Set LLC brice@emptysetllc.com emptysetllc.com