THE COALITION FOR COMMON SENSE IN GOVERNMENT PROCUREMENT

Jeffrey Calhoun
Supervisory Contract Specialist
General Services Administration
1800 F Street, NW Washington, DC 20006
Submitted via maspmo@gsa.gov

Subject: Comments on proposed AI Clause, GSAR 552.239-7001

Dear Mr. Calhoun:

The Coalition for Government Procurement ("Coalition") appreciates the opportunity to comment on GSA's proposed AI Clause, GSAR 552.239-7001.

By way of background, the Coalition is a non-profit association of firms selling commercial services and products to the Federal Government. Its members collectively account for a significant percentage of the sales generated through GSA contracts, including the MAS program. Coalition members also are responsible for many of the commercial item solutions purchased annually by the Federal Government. These members include small, medium, and large business concerns. The Coalition is proud to have collaborated with Government officials for 45 years in promoting the mutual goal of common-sense acquisition.

The Coalition supports the underlying objective of the proposed Government Terms and Conditions for AI Systems clause (the "AI clause") to safeguard Artificial Intelligence ("AI") systems and protect Government data. However, as discussed more fully below, Coalition members have concerns regarding the overly broad nature of the AI clause and have expressed that the AI clause may have a chilling effect on contractor use and adoption of AI should the clause be finalized as written. While the Government's intent to protect sensitive data and ensure AI system integrity is understandable, the implementation of these requirements raises significant concerns about practicality, enforceability, and market accessibility. Unnecessarily burdensome requirements can drive companies out of the Government marketplace, thus hampering agency access to innovation needed to meet their missions and leaving Government less, not more, secure.

Members of the Coalition ("members") have expressed feedback on the following topics:

- Applicability
- Definitions
- Service Provider Responsibility
- Intellectual Property Rights
- Disclosures and Documentation
- Incident Reporting

- Unbiased AI Principles and Government Evaluations
- Timing and Implementation

Applicability and Definitions

Applicability

The prescription at GSAR Part 539.71 states the Contracting Officer must insert the AI Clause into all solicitations and contracts for "Artificial Intelligence capabilities," yet the rule does not define "Artificial Intelligence capabilities." It appears GSA intends to incorporate the AI clause into all GSA MAS contracts, even where inclusion of the clause does not accomplish the Government's stated goals with respect to regulation of AI. This creates a serious risk of overbreadth and unintended consequences.

Recent regulatory experience underscores the danger of overly expansive AI definitions. For example, during California Privacy Rights Act ("CPRA") rulemaking, California faced significant criticism that an expansive definition would sweep in tools that merely assist human decisions—capturing a wide range of non-AI technologies, such as, spreadsheets, calculators, routine analytics, and creating unworkable compliance burdens. California ultimately narrowed the approach to focus on technologies that "replace" or "substantially replace" human decision-making, precisely to avoid this kind of accidental overinclusion and to preserve a workable, risk-based scope.

Further, the text of the AI clause appears to expand its application to contractors not only providing but also "using as part of performance of this contract an AI System" which could reach incidental use (embedded commercial tools) and supporting use (custom workflows that aid performance but do not drive deliverables). In doing so, the AI clause would place undue and impractical burdens on MAS holders, shrinking the field of willing vendors and increasing prices due to reduced competition.

- **Recommendation:** Provide a clear definition of "Artificial Intelligence capabilities" and limit the clause to GSA MAS contracts where the contractor is providing AI systems as a priced GSA Schedule product, rather than applying the AI clause to any MAS offering that may use AI as a component of a commercial product. A reasonable approach might categorize AI usage into tiers such as incidental use (embedded commercial tools), supporting use (custom workflows that aid but do not drive deliverables), and core use (AI systems that directly produce contract outputs or make decisions). Incidental use and supporting use cases should be explicitly out of scope of this clause. The Government could also consider incorporating the AI clause at the Order level on a case-by-case basis where the scope of work is explicitly for the purpose of providing the defined "AI capabilities."

Defining "American AI" and Non-US Entity AI Components

The current definition requires that an AI system be "developed and produced in the United States" and the AI clause prohibits the use of any AI components "manufactured, developed, or controlled by non-U.S. entities." Proposed GSAR 552.239-7001(a); (e)(2). This latter restriction is overly broad and creates practical barriers for any modern AI product in today's global technology landscape. Modern AI systems are built on complex, interconnected supply chains

that often involve international collaboration. Even companies headquartered in the U.S. frequently utilize global cloud infrastructure, incorporate training data from international sources, employ developers located overseas, and rely on open-source components developed by international communities. This restriction causes issues with how global companies leveraging a global workforce with local subsidiaries may develop products. It would effectively mean companies cannot sell a product to the U.S. Government that includes open-source models, where there is limited to no ability to verify where the open-source models were developed. The effect is to eliminate broad classes of commercially and technologically relevant AI components from eligibility, even where all processing of Government data occurs within a controlled, compliant environment.

Relatedly, if GSA implements de facto data localization requirements by requiring exclusively U.S.-origin development and production and prohibiting non-U.S. entity components, it may detrimentally impact the current Administration's efforts to persuade foreign allies (including the EU) to loosen their own data localization requirements. A U.S. procurement mandate that effectively localizes AI development, components, or data handling could be cited by foreign regulators as justification for maintaining or expanding localization policies that restrict U.S. exporters.

The clause does not provide a clear definition of what constitutes "developed and produced in the United States," creating ambiguity about whether a system qualifies under the following circumstances:

- its supporting workforce is equally distributed between the U.S. and another country;
- it has any international components in its supply chain;
- the entity supporting the system is not headquartered in the U.S., but has substantial business there; and/or
- the entity supporting the system is headquartered in the U.S. but the LLM is located outside the U.S.

Without further clarification, a broad reading of this requirement could exclude many high-quality, commercially available AI solutions from government consideration, significantly limiting competition and potentially increasing costs while reducing innovation. Furthermore, it may be virtually impossible for contractors to guarantee that every component, from semiconductors to training data sources, has no foreign elements. The 30-day post-award certification deadline is operationally impossible for complex AI stacks. Proposed GSAR 552.239-7001(e)(1).

- **Recommendation**. GSA should provide clear guidance on what "produced and developed in the United States" means in the context of modern AI supply chains, including treatment of open-source components, training infrastructure location, and corporate nationality of contributors. A practical standard might focus on the entity that controls the model and where the model weights are hosted and operated, rather than trying to trace the national origin of every component. This could incorporate existing product country of origin tests applied to AI Systems. This would allow the U.S.

Government to benefit from innovative AI tools without unnecessarily eliminating certain AI Systems.

"Government Data" and "Data Outputs"

The AI clause defines "Government Data" as "Data Inputs and Data Outputs." This seemingly provides government ownership over all inputs *and* all outputs (including metadata, logs, synthetic data, and other information), which is extraordinarily expansive. If a contractor uses an AI coding assistant and the AI generates code, per the clause it appears the Government owns that output even if the AI was simply accelerating work the contractor would have done anyway. Furthermore, the definition of "Data Outputs" includes "any. . .action produced by the AI System." Proposed GSAR 552.239-7001(a). The definition is so expansive that following it to its logical conclusion would require saving every interim prompt response and considering that to be government-owned data. This is not only impractical, but it would prove overwhelming for contractors attempting to comply, discouraging innovative providers from participating in GSA procurements.

In addition, many companies' AI enabled products generate recommended actions drawn from a finite, predefined set of permissible actions within the product's functionality. For example, certain systems may provide outputs as a result of simple functional system commands that are not creative or unique outputs generated from Government data. Under the current definition, the Government would obtain ownership rights not only in the outputs generated in response to Government prompts but also in the underlying action taxonomy itself. That would be far too broad and would effectively transfer ownership of functional components of contractor products.

Additionally, the definition of Data Outputs should include a background intellectual property ("IP") exclusion. Many companies' systems perform Retrieval-Augmented Generation ("RAG") or summarization over product documentation to assist users in troubleshooting. The Government's ownership of Data Outputs should not extend to underlying proprietary documentation or other background IP that may appear within or be referenced by an output.

- **Recommendation**: "Government Data" should be defined more narrowly to cover data that the Government provides to the contractor and deliverables produced under the contract, not every incidental AI interaction. Or, Government ownership of output should be limited to actual output from the system, not actions. It should include an exclusion for background IP from output ownership. Absent these revisions, the rule would deter commercial AI providers from offering leading products to the Government.

Service Provider Responsibility

Section (c) of the clause states, "the Contractor is responsible for the Service Provider's adherence to this clause." GSAR 552.239-7001(c). The definition of "Service Provider" in Section (a) is broad, to include any "entity that directly or indirectly provides, operates, or licenses an AI system but is not a party to the contract." GSAR 552.239-7001(a). Contractors will have difficulty ensuring full compliance by Service Providers (particularly indirect providers) where commercial LLM/API providers often have standard terms that may conflict with certain aspects of the clause. These commercial providers may bring innovative and

effective solutions that the Government will not be able to use based on the overly restrictive nature of the clause.

- **Recommendation**: Revise the definition of "Service Provider" to mean only direct providers and/or subcontractors selected by the contractor. Update the clause to state specific requirements for contractors to ensure Service Providers meet (similar to the stated requirements for external Cloud Service Providers in DFARS 252.204-7012(b)(2)(ii)(D)) so that the diligence burden on contractors is reduced and contractors can focus on ensuring compliance by Service Providers with the stated requirements rather than the entirety of the clause.

Intellectual Property Rights

Ownership and Restrictions on Use

Members have concerns regarding the potential impact of GSAR 552.239-7001(d) – Intellectual Property Rights, particularly as it relates to the ownership and reuse of AI-related developments created during contract performance.

Section (d) appears to require the contractor to relinquish ownership of certain custom developments and Government data created under the contract. From a contractor perspective, this requirement may create unintended risk by limiting the contractor's ability to reuse internally developed methodologies, tools, models, or enhancements in future contracts or commercial applications.

The clause also raises potential concerns related to commingling of IP and funding sources. In practice, AI development frequently involves the use of existing contractor-owned algorithms, frameworks, training methods, and datasets that may be refined or expanded during contract performance. Without clear boundaries between contractor pre-existing IP, contractor-funded enhancements, and Government-funded developments, there is a risk that contractor-developed capabilities could unintentionally become subject to Government ownership or broader rights than intended.

To mitigate this risk, contractors would need to implement strict version control, configuration management, and cost segregation practices to clearly demonstrate what was developed using Government funding versus contractor or commercial funding sources. These are not marginal compliance steps; they are expensive, ongoing operational burdens that will discourage contractors from incorporating their most innovative proprietary capabilities into Government solutions.

Critically, these IP and documentation burdens will predictably shrink the pool of contractors willing and able to compete for covered procurements. When fewer commercial innovators can participate — particularly those whose value depends on protecting and reusing proprietary AI capabilities – the Government's access to best-in-class solutions diminishes, competition declines, and prices rise. In short, the proposed IP approach risks turning commercial AI procurement into tailored, government-unique development: slower to deliver, costlier to maintain, and less competitive over time.

- **Recommendation:** The AI clause should provide clearer distinctions between contractor background IP and Government-funded deliverables, as well as guidance on how derivative improvements or mixed-funding developments should be treated. Clarification in these areas would help ensure both protection of Government interests and preservation of contractor innovation incentives.

Model Output and "Discretionary Policies"

The AI clause states that an AI System must not refuse to produce outputs or conduct analysis based on the Contractor's or Service Provider's "discretionary policies." Proposed GSAR 552.239-7001(d)(2)(ii). This concept and restriction should be better defined. While we appreciate the Administration's underlying concerns, all modern foundation models contain embedded guardrails that prevent the model from generating harmful, dangerous, or illegal content (*e.g.*, instructions for weapons construction or encouraging self-harm). If these guardrails are considered "discretionary policies," then compliant models would need to remove model level- safety features. Although the AI clause notes that this does not require retraining or altering model weights, this requirement is unclear, appears inconsistent with industry standard- safety practices, and raises questions around how this restriction is supposed to be implemented. Are future models supposed to be trained in ways that have any post-training guardrails removed? Or will the Government just not purchase systems with models with guardrail capabilities? If so, the consequence is that the Government will not be able to procure current safe, commercially available models, further restricting potential U.S. Government capabilities and innovation.

- **Recommendation**. The AI clause should specify this prohibition is not intended to prohibit all refusal behaviors, such as legal, security, and essential safety safeguards.

"Eyes Off" Data Handling Requirements

Proposed GSAR 552.239-7001(d)(4)(ii) includes a requirement for "eyes off" data handing procedures. The requirement for "eyes off" data handling procedures that severely restrict human review of Government data is well-intentioned, but potentially unworkable in practice. While protecting sensitive government information is critical, the mandate that human access be limited to what is "strictly necessary" and that all such access be logged and justified creates significant operational challenges. Modern AI systems require human oversight for quality assurance, safety interventions when systems produce problematic outputs, technical debugging when systems malfunction, and continuous improvement processes. Customer support—a fundamental service expectation—becomes extraordinarily difficult when support personnel cannot examine the actual interactions that led to user issues. The requirement essentially asks contractors to provide effective AI services while being largely blind to how those services are actually performing with Government data. This creates a paradox where contractors are held responsible for system performance and compliance but are simultaneously prohibited from the kind of monitoring and intervention that would allow them to ensure quality. The logging requirements, while reasonable in principle, add substantial compliance overhead and may generate massive audit trails that become difficult to manage and review effectively. Furthermore, it is unlikely that contractors will be able to compel and enforce this regulation on commercial AI Service Providers.

- **Recommendation**: The Government should seek input from commercial AI Service Providers on this particular aspect of the clause and revise it to reflect what is practical and reasonable.

Restrictive Custom Development Terms

The AI clause's treatment of custom developments creates a significant disincentive for innovation during contract performance. *See* Proposed GSAR 552.239-7001(d)(5). Any modifications, customizations, configurations, or enhancements to AI systems developed specifically for the Government become the exclusive property of the Government in perpetuity. This is inconsistent with the standard government data rights regime, which generally provides the Government with unlimited rights to data first developed in performance of a contract, while the contractor retains ownership of IP developed. As written, the AI clause appears to dictate that even broadly applicable technical improvements, general performance optimizations, or bug fixes discovered during government work cannot be applied to the contractor's commercial products or used to benefit other customers. A contractor who develops an innovative approach to improving AI safety or accuracy while working on a government contract would be prohibited from applying that same innovation to their commercial offerings, even if the innovation has no specific connection to government data or requirements. This makes government contracts significantly less valuable from a business development perspective and may discourage contractors from investing in innovative solutions, knowing that any breakthroughs will be locked away for exclusive government use rather than advancing their commercial capabilities.

- **Recommendation**: Revise the IP Rights text to better align with the standard government data rights regime in the FAR (*e.g.*, FAR 52.227-14). The AI clause should be rewritten to strike the right balance between protecting the Government's right to own its data (namely, data inputs) but allow the contractor to retain rights to improve its model for the benefit of all customers.

Contractor Obligations (Section (e))

Disclosures & Documentation

The AI clause states the contractor "must disclose all AI systems used in performance of this contract" and disclose "whether the AI system has been modified or configured to comply with any non-U.S. federal government or commercial compliance or regulatory framework…" Proposed GSAR 552.239-7001(e)(1). This disclosure requirement is broad and requires more specificity for contractors to comply.

Further, the AI clause requires documentation on "AI System decision-making processes, logic, and operational parameters," "Testing methodologies used to detect and mitigate noncompliance with unbiased AI principles described in paragraph (i)(1)," as well as other information. Proposed GSAR 552.239-7001(e)(6). AI systems (especially Large Language Models ("LLMs")) are probabilistic and indeterminate. That means the contractor cannot document the exact internal reasoning for every output the way it can with traditional rule-based software. To the extent contractors are using third-party commercial AI Service Providers, it may not be feasible for a contractor to impose these documentation requirements. Commercial model Service

Providers generally do not disclose model internals or proprietary testing frameworks, and contractors cannot compel them to do so.

- **Recommendation**: The clause should be updated to clarify the disclosure requirement applies to AI Systems that directly produce contract deliverables, process Government Data, or make decisions material to contract performance. AI-enabled commercial tools used incidentally in the course of contract performance should be explicitly out of scope. The documentation requirement should be restricted to direct providers of, or resellers for, the AI systems being delivered to the Government—not contractors that merely use AI tools incidentally to support performance. The clause should also acknowledge that many LLMs and AI systems cannot reliably document internal reasoning or generate complete, step-by-step explanations. Accordingly, we propose a carve-out for requirements such as mandatory summaries of intermediate processing steps, model routing rationales, and similar observability measures that may not cleanly map onto how many current AI systems actually function. The AI clause should be revised to acknowledge that LLMs and AI systems may not be able to document all internal reasoning.

Incident Reporting

The proposed clause includes a section on incident reporting, requiring the contractor or Service Provider to submit a CISA incident reporting form, notify the Government of an incident within 72 hours, and provide daily status updates until the incident is resolved. *See* Proposed GSAR 552.239-7001(e)(4). This incident reporting requirement is duplicative of other reporting requirements already included in government contracts and adopts a stricter regime than the EU AI Act. Contractors already are required to notify the Government of incidents in which sensitive or controlled government information is impacted. Further, as noted in the proposed clause, AI Systems with FedRAMP authorization will already be required to follow incident notification procedures dictated through FedRAMP. Further, CISA's forthcoming incident reporting for critical infrastructure regulations will also require reporting of incidents within 72 hours. In light of the federal Government's stated goal to harmonize federal cybersecurity and incident reporting regulations, the proposed GSAR clause should not introduce yet another incident reporting provision that will be duplicative of existing and forthcoming regulations without significantly advancing the Government's mission.

- **Recommendation:** The incident reporting subsection should be removed from the AI clause. To the extent the incident reporting subsection is not removed, it should be updated to specify an incident "impacting AI systems used in performance of this contract." Further, requiring daily status updates is overly broad and onerous. This should be updated to "periodic updates" as the contractor and/or Service Provider have material updates as they investigate the incident.

Unbiased AI Principles and Government Evaluations

The AI clause incorporates prohibitions against models producing outputs that are "ideological," "non-neutral," or influenced by concepts such as diversity, equity, and inclusion ("DEI"). Proposed GSAR 552.239-7001(i)(1). It requires companies to use "commercial efforts" to ensure truthfulness, neutrality, and non-partisanship. The mandate that AI systems be "truthful" and

prioritize "historical accuracy, scientific inquiry, and objectivity" sounds reasonable in the abstract, but becomes problematic in application. Many historical and scientific questions involve contested interpretations, evolving understanding, and legitimate scholarly disagreement. Requiring an AI system to provide "the truth" on complex, multifaceted topics is asking for something that may not exist in singular form.

The Government reserves the right to conduct automated assessments using its own benchmarks, methodologies, and testing criteria at any time, while simultaneously classifying all such benchmarks and methodologies as protected Government data that need not be shared with contractors. *See* Proposed GSAR 552.239-7001(i)(2). Unlike transparent compliance frameworks where requirements are clearly specified and testable, this approach creates a "black box" evaluation system where contractors can be found non-compliant based on criteria they were never given the opportunity to address.

Contractors face serious consequences and penalties based on Government evaluation results, including suspension or termination and decommissioning costs. Government agencies might integrate AI systems deeply into their operations, creating extensive dependencies and making migration to alternative systems expensive and time-consuming. Since contractors have little control over how extensively agencies choose to rely on their systems, they could face enormous decommissioning costs that bear no relationship to the revenue generated by the contract. By introducing subjective ideological-bias standards and pairing them with undisclosed government benchmarks and potential decommissioning liability, the clause creates profound uncertainty for contractors, who may reasonably fear that good-faith disputes over "ideology" or "neutrality" could become a basis to exit contracts early and shift migration costs to the contractor. This chilling effect will reduce the number of companies willing to contract with the U.S. Government, diminishing competition and innovation and ultimately harming the Government's access to best-in-class AI solutions. The combination of subjective compliance standards and unlimited financial liability creates an untenable risk profile that may discourage qualified contractors from bidding on government AI contracts altogether.

- **Recommendation**. Eliminate this requirement altogether. Alternatively, update the neutrality requirements and benchmarks to align with objective, published standards. Contractors should have access to the evaluation criteria so they can proactively ensure compliance. The clause should also explicitly carve out safety-critical content moderation (*e.g.*, preventing generation of dangerous or illegal content) from the prohibition on output refusals. The clause should include a process for the Government to provide the contractor with reasonable notice and an opportunity to cure should an evaluation issue be identified.

Other Considerations

The "Nature of AI" Is Moving Faster Than This Clause

The reality of how AI products work in 2026 is inconsistent with several assumptions that appear to be included in the AI clause. For example, the change management provision requires 30 days' notice before "material changes." Cloud-based AI providers update models continuously. A contractor using an Application Programming Interface ("API") based AI service does not

control (and often does not even know about) incremental model changes. Requiring 30 days' notice for something the contractor cannot predict or control is structurally unworkable. The seven-day bias notification requirement has the same problem – a contractor may not be able to detect that a routine model update "materially increased output bias" without running continuous benchmarking.

Agentic AI blurs the lines further. The AI clause requires summarizing "intermediate processing actions and decision points" for systems using agentic processes or RAG. But agentic AI systems can chain dozens of tool calls, spawn sub-agents, and make routing decisions dynamically. Full traceability is a worthy goal, but the clause does not acknowledge that the level of observability varies dramatically by platform and architecture. Some of this instrumentation simply does not exist yet for commercial AI products.

AI-as-infrastructure is becoming invisible. We are rapidly moving toward a world where AI is not a discrete "system" you can point to clearly; rather, it is a layer running beneath databases, analytics platforms, cloud services, and operating systems. Requiring contractors to identify and disclose every instance of AI touching contract performance will become increasingly impractical as AI becomes ambient infrastructure.

These practical realities underscore the need for a faster, simpler, and more flexible procurement framework — consistent with the December 11, 2025 Executive Order, "Ensuring a National Policy Framework for Artificial Intelligence," which emphasizes speed and ease of adoption and recognizes that AI is developing too quickly for rigid, fragmented regulation (including by proposing federal preemption of state AI laws). GSA should align the clause accordingly by adopting outcome-based, risk-tiered requirements rather than prescriptive operational mandates.

- **Recommendation:** The clause should focus on *outcomes and risk* rather than trying to regulate the technology itself. Require contractors to maintain a risk-aware posture and disclose AI use that materially affects deliverables or decision-making, rather than mandating exhaustive cataloguing of every AI touchpoint. The clause should also acknowledge the shared-responsibility model of cloud AI. Contractors should not be held responsible for aspects of a Service Provider's platform they have no visibility into or control over.

<u>Harmonization</u>

The Administration has explicitly prioritized reducing regulatory burden and harmonizing compliance frameworks across agencies, particularly for commercial providers. The AI clause creates a new, GSA-specific AI compliance regime that goes beyond the stated objectives of OMB M-25-21, OMB M-25-22, or existing governmentwide AI policy framework. Rather than harmonizing with existing rules, it layers additional obligations on top of FedRAMP, FISMA, and NIST AI RMF — increasing cost and complexity without a corresponding security or performance benefit.

Furthermore, the AI clause as written seems to include elements that are already covered by long-standing existing clauses in government contracts, including data rights and cyber incident reporting. Duplication in overlapping clauses creates ambiguity and complexity for contractors and the Government alike.

- **Recommendation:** Eliminate elements in this clause that are adequately addressed by existing clauses.

Timing and Implementation

Given the complexity, novelty and importance of this clause and the volume of comments, we strongly recommend that implementation of GSAR 552.239-7001 go through the formal rulemaking process, with an opportunity for notice and comment, to provide industry sufficient time to address the substantial obligations imposed on contractors and service providers and ensure that the Government has access to this evolving, cutting-edge technology from the widest pool of vendors. Contractors will need time to review current AI Provider commercial agreements for conflicts with the clause and for active subcontractors to address impacts. Thus, the Coalition recommends GSA further engage with industry to develop a reasonable timeline for implementation to include formal rulemaking.

Conclusion

The Coalition hopes you find these comments useful and thanks you for your time and consideration. Should you have any questions or concerns, I may be contacted at rwaldron@thecgp.org or 202-331-0975.

Sincerely,

Roger Waldron
President
The Coalition for Government Procurement