Submitted via https://www.regulations.gov. on 8/3/2026

Subject: Notice-MVAC-2026-01 - Comments on Proposed GSA Clause, GSAR 552.239-7001, *Basic Safeguarding of Data within Large Language Model Artificial Intelligence Systems (LLMs)*

The Coalition for Common Sense in Government Procurement (Coalition) appreciates the opportunity to comment on the General Service Administration's (GSA's) proposed Artificial Intelligence (AI) clause, General Services Administration Regulation (GSAR) 552.239-7001 (AI clause).

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 Multiple Award Schedule (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 submitted extensive comments with respect to the initial draft AI clause that was proposed as part of MAS Refresh 31. We would like to commend GSA for accepting and considering industry input and deciding to take more time and receive additional input on this important, complex, cutting-edge topic. Collaboration between industry and Government is the best way to arrive at a solution that meets the needs of the Government without potentially depriving the Government of the best that the commercial market has to offer.

However, as discussed more fully below, Coalition members continue to have concerns regarding the overly broad language utilized in the AI clause and believe 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 preservation of a contractor's background intellectual property. Members also have concerns about the practicality of contractor implementation of the burdensome requirements. For example, several requirements in the AI clause assume that a contractor owns and operates the underlying Large Language Model (LLM). Most

contractors, however, operate as an LLM System Integrator or LLM User. Contractors can only be responsible for matters within their control, and contractors operating in these roles cannot as a general matter guarantee model neutrality, prevent commercial providers from conducting human content review, or compel the level of intermediate-step disclosure the proposed clause requires.

In its June 17, 2026 request for comments, GSA requested input from the public to understand the scope and impact of the AI clause with respect to five questions. Coalition members have the following concerns or input with respect to each of these questions.

***

1. Does the change in clause prescription adequately address previous concerns about the broadness of the scope of the clause?

While the revised AI clause addresses certain of the Coalition's previously expressed concerns regarding the breadth of its scope, the clause as currently drafted continues to depart from common commercial practices and is inconsistent with the Administration's stated objective of promoting commercial AI adoption. Specifically, the following concerns remain:

<table>
<tr><th colspan="2">Comment #1</th></tr>
<tr><td colspan="2"><u>Reference</u>: Part 539.71 (a), Applicability

<u>Language at issue</u>: "The contracting officer must insert the clause at 552.239-7001, Basic Safeguarding of Data Within Large Language Model Artificial Intelligence Systems, in solicitations and contracts, including those for commercial products and services, when Government data will be <u>processed by</u> a LLM." (emphasis added).</td></tr>
<tr><td><u>Description of Issue</u>: The applicability trigger is ambiguous as to what constitutes "processed by" a LLM and risks overbroad, precautionary inclusion of the clause in contracts.

The clause prescription states that this clause applies only when "Government data will be processed by a LLM." This language is ambiguous as to what constitutes "processed by" an LLM. Modern LLM-based systems frequently involve multi-stage pipelines in which Government Data may be parsed, cleaned, or otherwise pre-processed by ancillary components before being submitted to the LLM itself. Neither the prescription nor the clause specify whether such upstream or downstream processing steps fall within the scope of "processed by" the LLM, or whether the trigger is limited to data that is directly ingested by the model's inference engine. This ambiguity is compounded by the risk that Contracting Officers will include the clause in contracts — particularly indefinite-delivery/indefinite-quantity (IDIQ) vehicles – on a precautionary basis, even where the applicability of the clause to a given task order is uncertain. Given the substantial compliance, reporting, and</td><td><u>Proposed Clarification or Recommendation</u>: The term "processed by a LLM" should include only data directly submitted to the model's inference engine. It should not capture upstream and downstream processing by ancillary system components. Additionally, we recommend that GSA require contracting officers to produce documented, affirmative determinations that Government Data will be processed by an LLM before the contracting officers can include the clause in the resulting contract.</td></tr>
</table>

documentation obligations imposed by paragraph (f), overbroad or precautionary inclusion of the clause would impose significant and unwarranted burdens on both contractors and the Government. This potentially overly broad application is also contrary to the Administration's intent via memorandum M-25-22, which states that the government "must communicate clear and specific requirements that make it easy for vendors to offer state-of-the-art AI capabilities to support efficient and effective public services."	

Comment #2	
Reference: Part 539.71(b)(1), Applicability ***Language at issue:*** "The LLM is **embedded in** a **common commercial product**, such as a word processor or map navigation system." (emphasis added)	
Description of Issue: ***The "embedded in a common commercial product" exception is too narrow to capture the full range of technical architectures through which LLM functionality is integrated into commercial products, and the term "common commercial product" is undefined.*** The prescription exception is limited to instances where the LLM is "embedded in" a "common commercial product." We have two separate concerns with this exception. First, the term "common commercial product" is undefined and its relationship to the Federal Acquisition Regulation (FAR) 2.101 definition of "commercial product" is unclear. The illustrative examples provided (word processors and map navigation systems) suggest consumer-grade tools, raising the question of whether enterprise software platforms with integrated AI copilot functionality qualify for the exception. And, as new tools become more mainstream, there is a question regarding at what point a commercial product may be considered "common." Second, this language, particularly use of the term "embedded in," is unnecessarily narrow. In practice, LLM functionality is integrated into commercial products through multiple technical architectures, including direct embedding, application programming interface (API) calls, and third-party integrations. And the term "embedded" does not capture these alternative configurations. A contractor using a commercial product that accesses LLM functionality via API rather than through an embedded model would fall outside this exception despite presenting the same risk profile.	***Proposed Clarification or Recommendation***: We recommend that GSA replace "embedded in" with "used with, integrated into, or embedded in" to capture the full range of technical architectures, and replace "common commercial product" with "commercial product" as defined in FAR 2.101. GSA should also clarify that the exception extends to enterprise software platforms with integrated AI capabilities, not solely consumer-grade tools. AI/LLM features embedded in commercial products should be excluded from the scope of this clause.

Comment #3

Reference: Part 539.71(b)(2), Applicability ***Language at issue:*** "The LLM functionality is **incidental to the primary purpose** of the core requirement being procured." (emphasis added).	
Description of Issue: ***The phrase "incidental to the primary purpose" lacks objective criteria, risking overbroad application of the clause to routine internal use of commercial AI tools.*** The term "incidental" as applied to LLM functionality is undefined, and the related concept of "primary purpose of the core requirement being procured" lacks objective criteria for determination. In the absence of a definition or bright-line test, Contracting Officers are likely to default to broad application of the clause, capturing routine internal use of commercial AI tools by contractor employees for drafting, summarizing, researching, and code assistance on contracts that are not AI procurements. Where the contract deliverables do not require delivery of an LLM or AI-enabled capability, and the contractor merely uses an LLM as an internal productivity tool to produce deliverables, the clause's obligations (including license grants, data handling requirements, and flowdown provisions) are disproportionate and inapplicable to the actual risk.	***Proposed Clarification or Recommendation***: GSA should (1) define "incidental" LLM functionality with a bright-line test to exclude AI tools used in the course of producing a professional services deliverable, where the contract does not require delivery of an AI-enabled capability, and no Government-specified AI tool is mandated; and (2) state that the clause applies at the contract level, task order level, or to segregable contract line items.

Comment #4	
Reference: 552.239-7001 Paragraph (c), Order of Precedence ***Language at issue***: "For purposes of the order of precedence in GSAR clause 552.212-4, this clause is incorporated into the 'schedule of supplies/services'. This clause establishes specific requirements that take precedence over conflicting provisions in the Contractor's policies, requirements, terms, conditions, or commercial agreements."	
Description of Issue: ***Incorporating the AI clause into the highest-precedence tier of the contract order of precedence overrides all commercial contract terms — including the Contractor's negotiated agreements with commercial LLM providers – and imposes compliance obligations that the Contractor may have no contractual authority to fulfill.*** Under GSAR 552.212-4, the "schedule of supplies/services" occupies the highest position in the order of precedence — above other commercial Federal Acquisition Regulation (FAR) clauses, addenda (including any Contractor Agreement or commercial terms and conditions), solicitation provisions, and other attachments. By incorporating the AI clause into this highest-precedence tier, GSA is placing its AI safeguarding requirements above all other commercial contract terms, including the Contractor's standard commercial agreements and any negotiated terms with its	***Proposed Clarification or Recommendation***: We recommend that GSA revise paragraph (c) to limit the AI clause's precedence override to provisions directly governing the security, integrity, and confidentiality of Government Data, rather than granting blanket precedence over all commercial contract terms. GSA should also incorporate: (1) a safe harbor provision establishing that the Contractor shall not be held liable for noncompliance

commercial LLM providers. This creates two distinct concerns. First, the Contractor cannot be held liable for compliance obligations it has no contractual authority to flow down to commercial vendors whose standard terms of service preclude such obligations. The AI clause imposes numerous requirements — including data handling, audit logging, human access restrictions, and bias monitoring – that depend on the cooperation or configuration of the underlying LLM provider. Where the Contractor operates as an LLM User or System Integrator rather than the LLM Developer, the Contractor lacks the technical access or contractual leverage to compel the commercial provider's compliance. Second, the clause requires the Contractor to certify or provide information that commercial LLM providers may not disclose. For example, paragraph (f)(3) requires the Contractor to certify within 30 days of award whether the LLM has been configured to comply with non-U.S. statutes or policies — information that commercial LLM providers typically do not share with downstream customers. Placing the AI clause at the highest precedence tier does not give the Contractor the ability to obtain such information; it merely exposes the Contractor to noncompliance liability for circumstances beyond its control.	with any requirement of the AI clause where such noncompliance is attributable to a commercial LLM provider's refusal to modify its standard terms of service, disclose proprietary information, or otherwise cooperate with the Contractor's compliance efforts; and (2) a commercially reasonable efforts standard for any obligation that requires the Contractor to obtain information, certifications, or commitments from a commercial LLM provider, in recognition that the Contractor's ability to compel such cooperation is limited by the commercial provider's standard terms of service.

2. Are the requirements such as Government data ownership and protection and contractor accountability clearly defined?

Coalition members have raised several concerns regarding the definition and scope of Government data ownership and contractor accountability. In particular:

A. Data Ownership and Protection

Comment #1	
Reference: 552.239-7001 Paragraph (b), Definitions ***Language at issue:*** "*Background Data* means any **pre-existing** proprietary content, reference materials, knowledge bases, or other intellectual property owned or controlled by the Contractor that may be referenced, retrieved, augmented, or otherwise incorporated into the LLM's processing or outputs through any enrichment mechanism, including but not limited to retrieval processes, vector stores, embeddings, knowledge graphs, plugins, tool calls, agent actions, or similar mechanisms." (emphasis added).	
Description of Issue: ***The definition of "Background Data" fails to cover Contractor proprietary data developed independently and at private expense after contract award but outside the contract's scope.***	***Proposed Clarification or Recommendation:*** We recommend that GSA revise the definition of "Background Data" to include not only pre-existing

The definition of "Background Data" is limited to "pre-existing proprietary content" owned or controlled by the Contractor. Under established federal procurement data rights frameworks (e.g., FAR 52.227-14, Defense Federal Acquisition Regulation Supplement (DFARS) 252.227-7013), the analogous concept of "limited rights data" or "restricted computer software" encompasses not only data existing prior to the contract but also data developed by the contractor at private expense outside the contract. The current definition of Background Data omits this second category — data developed by the contractor independently and outside the scope of the contract but after contract award – thereby leaving such data without a clear classification under the clause and potentially exposing it to treatment as Government Data or Data Outputs.	proprietary content but also proprietary content developed by the Contractor independently and at private expense outside the scope of the contract, consistent with the treatment of contractor-developed data under FAR 52.227-14 and DFARS 252.227-7013. This revision would ensure that contractor intellectual property developed after contract award but outside the contract is appropriately classified and protected.

Comment #2	
Reference: 552.239-7001 Paragraph (b), Definitions ***Language at issue:*** "*Data Inputs* means all data, information, personally identifiable information (PII), or content submitted to the LLM and **related operational systems** by, or **created for, the Government**, including but not limited to user prompts, queries, instructions, system prompts, source data, documents, knowledge bases, Government email addresses, user account information, and any other information or content submitted to the LLM and related operational systems by or on behalf of the Government. This definition specifically excludes any background intellectual property or information existing prior to entry into this contract or task/delivery order or developed independently by such Contractor." (emphasis added).	
Description of Issue: ***The definition of "Data Inputs" is overbroad due to the undefined term "related operational systems" and unclear interaction with the Background Data exclusion.*** The definition of "Data Inputs" is overbroad in two respects. First, it captures all data submitted to "the LLM and related operational systems," but the term "related operational systems" is undefined and could encompass any system in the contractor's technology stack that interacts with the LLM. Second, the definition includes data "created for" the Government, which could sweep in contractor proprietary data that is incorporated into a deliverable or used in a database that also serves Government purposes. Although the definition contains an exclusion for "background intellectual property or information existing prior to entry into this contract or task/delivery order or developed independently by such Contractor," the interaction between this exclusion and the scope of the definition is unclear. Specifically, if excluded Background Data is submitted to the LLM as an	***Proposed Clarification or Recommendation***: We request that GSA define or narrow the term "related operational systems" and clarify that Background Data retains its excluded status even when submitted to the LLM as an input. The definition should be narrowed to focus on substantive content originating from the Government, rather than capturing all data submitted to the system by or on behalf of the Government.

input, the clause does not address whether that data retains its excluded status or is reclassified as a Data Input subject to Government ownership and use restrictions.	

Comment #3	
Reference: 552.239-7001 Paragraph (b), Definitions ***Language at issue:*** "*Data Outputs* means all data, information, PII, any improvements, enhancements, corrections, annotations, or other modifications made to Data Inputs, or content generated by the LLM in the performance of this contract, including but not limited to responses, results, analyses, anonymized data, derivative data, metadata, logs, synthetic data, and any other output or action produced by the LLM, **regardless of whether such output incorporates or is derived from Background Data**. This definition specifically excludes technical system-level data that contains no Government information or Government usage context, such as performance metrics, token counts, and processing times." (emphasis added).	
Description of Issue: ***The definition of "Data Outputs" asserts Government ownership over LLM outputs regardless of whether they incorporate the Contractor's pre-existing Background Data, creating a significant risk to Contractor intellectual property rights.*** The definition of "Data Outputs" encompasses all content generated by the LLM in the performance of the contract, "regardless of whether such output incorporates or is derived from Background Data." This language, read in conjunction with the Government's assertion of full ownership over Data Outputs under paragraph (e)(1)(i), creates a significant risk to contractor intellectual property rights. Where an LLM output incorporates or is derived from the contractor's Background Data, the Government would acquire ownership of an output that is substantively derived from the contractor's pre-existing intellectual property. The definition requires a carve-out preserving the contractor's rights in any portion of a Data Output that constitutes or is derived from Background Data.	***Proposed Clarification or Recommendation***: We recommend that GSA (1) confirm that general-purpose methodologies and tools not specifically funded as a contract deliverable remain contractor-owned and (2) revise the definition of "Data Outputs" to include an express carve-out that the contractor retains all contractor intellectual property and privately funded improvements thereto containing no government data, including, *inter alia*, telemetry, security logs, usage metrics, platform metadata, schemas, model know-how incorporated into or reflected in Data Outputs. The Government's ownership interest should extend only to the novel content generated by the LLM using government funding for specific contract deliverables that does not constitute or derive from the contractor's Background Data.

Comment #4	
Reference: 552.239-7001 Paragraph (e)(1)(i), Rights in Government Data ***Language at issue:*** "The Government retains full ownership of, and will own, all Government Data and Custom Developments. The Contractor does not have any rights to use Government Data or Custom Developments provided to the Contractor other than those described in paragraph (e)(1)(iii)."	
Description of Issue: ***The Government's assertion of full ownership over all Government Data — which includes Data Outputs derived from Background Data — could effectively transfer ownership of the Contractor's pre-existing intellectual property to the Government without compensation.*** Paragraph (e)(1)(i) provides that "The Government retains full ownership of, and will own, all Government Data and Custom Developments." Because "Government Data" includes Data Outputs, and Data Outputs encompass content generated by the LLM "regardless of whether such output incorporates or is derived from Background Data," the Government's ownership assertion extends to outputs substantively derived from the contractor's pre-existing proprietary data — including general-purpose prompt engineering, Retrieval-Augmented Generation (RAG) architectures, and implementation methodologies. This creates a circular dynamic: the contractor's Background Data, once processed by the LLM and reflected in an output, is reclassified as Government-owned Data Output, effectively transferring ownership of the contractor's intellectual property to the Government without compensation or negotiation.	***Proposed Clarification or Recommendation***: We request that GSA add a savings clause to paragraph (e)(1)(i) providing that the Government's ownership of Government Data and Custom Developments does not extend to, and does not diminish the contractor's rights in, Background Data, even where such Background Data is incorporated into or reflected in Data Outputs. The Government's ownership interest should be limited to Government-originated content and novel LLM-generated content that does not constitute or derive from the contractor's pre-existing intellectual property.

Comment #5	
Reference: 552.239-7001 Paragraph (e)(1)(iv), Rights in Government Data ***Language at issue:*** "To the extent the Contractor obtains any intellectual property rights in Government Data, or **any improvements, enhancements, feedback, or derivative works thereof**, the Contractor assigns and transfers all such rights to the Government effective immediately upon creation." (emphasis added).	
Description of Issue: ***The automatic assignment provision operates to transfer the Contractor's intellectual property rights in any output incorporating pre-existing Background Data, which is inconsistent with established Federal data rights frameworks.***	***Proposed Clarification or Recommendation***: We recommend that GSA add an express limitation to

Paragraph (e)(1)(iv) provides that "to the extent the Contractor obtains any intellectual property rights in Government Data, or any improvements, enhancements, feedback, or derivative works thereof, the Contractor assigns and transfers all such rights to the Government effective immediately upon creation." The phrase "improvements, enhancements, feedback, or derivative works thereof" is dangerously broad. Read literally, it could pull in contractor-developed improvements to the underlying LLM that were merely *informed by* Government Data processing — even where the improvement is a general capability with broad commercial applicability. This is perpetuated by the fact that Paragraph (b) defines "Government Data" to include Data Outputs, and "Data Outputs" is defined to include content "derived from Background Data." This provision operates as an automatic assignment of the contractor's intellectual property rights in any output that incorporates the contractor's pre-existing proprietary content. This result is inconsistent with established Federal data rights frameworks, which recognize and protect contractor rights in data developed at private expense.	paragraph (e)(1)(iv) providing that the automatic assignment does not apply to, and does not transfer, any intellectual property rights in Background Data or in any portion of a Data Output that constitutes or is derived from Background Data. The assignment should be limited to intellectual property rights in Government-originated content and novel LLM-generated content that does not incorporate the contractor's pre-existing intellectual property.

Comment #6	
***Reference*:** 552.239-7001 Paragraph (e)(4)(ii), Government Data Handling and Processing Requirements. ***Language at issue*:** "The Contractor must implement Data Handling Procedures that **restrict human access to Government Data**. These protocols must be implemented at the contract level and customization at the individual contract or task/delivery order level. Automated processing systems and operational controls that prevent human personnel from the Contractor or any Third-party entity from viewing, accessing, or reviewing Government Data during normal operations include: (A) Automated data ingestion, processing, and response generation without human content review; (B) Technical access controls that prevent personnel from viewing Government Data; . . . (E) Audit logging systems that **track data processing activities** without capturing or displaying actual Government Data." (emphasis added).	
Description of Issue*:** ***The data handling requirements restricting human access to Government Data create an internal contradiction on professional services contracts, where the Contractor's authorized personnel must review LLM outputs to perform the work. Paragraph (e)(4)(ii) requires the Contractor to restrict human access to Government Data and implement automated controls that prevent Contractor and third-party personnel from viewing Government Data during normal operations. This requirement is drafted for backend infrastructure and LLM platform providers. On professional services contracts, however, the Contractor's authorized personnel must review, analyze, and act upon Data Outputs to perform the work — yet Data Outputs are a subset of Government Data, creating an internal contradiction. The clause fails to	***Proposed Clarification or Recommendation:*** We request that GSA clarify that paragraph (e)(4)(ii) governs the LLM Developer's and System Operator's backend personnel with system-level access and does not restrict the prime Contractor's authorized end users from viewing or acting upon Data Outputs in the course of contract performance. GSA should also address whether human review

distinguish between backend system-level access by the LLM Developer's or Operator's personnel and authorized end-user access by the prime Contractor's project team, and does not address whether human review for safety, quality assurance, or accuracy verification — including Government-directed review – is prohibited. Separately, paragraph (e)(4)(ii)(E) requires the Contractor to implement "[a]udit logging systems that track data processing activities." Practically speaking, for Contractors operating as LLM Users (where project team members use commercial AI tools as general productivity aids) this requirement would necessitate that every staff member log every individual use of an AI system, an obligation that is operationally unworkable at scale. The disclosure requirements under paragraph (f)(1), which already require identification of the LLM used in performance of the contract, are sufficient to document general LLM use by a project team without imposing granular, per-interaction logging obligations on individual end users.	for safety, abuse moderation, or quality assurance — including review directed by the Government – is prohibited or subject to specific conditions. Additionally, GSA should remove the audit logging requirement under paragraph (e)(4)(ii)(E) for Contractors operating as LLM Users, as the existing disclosure obligations under paragraph (f)(1) adequately document general LLM use by project teams and the per-interaction logging obligation is operationally unworkable for end-user personnel.

Comment #7	
***Reference*: *552.239*-7001 Paragraph (e)(4)(viii), Custom Developments and Model Rights** ***Language at issue*:** "If there is a requirement for custom developments, the Contractor must: (A) Dedicate the Custom Developments, including any customized or enhanced models resulting from such Custom Developments, to the Government's exclusive use; (B) Treat such Custom Developments, custom models, and all associated Data as the Government's confidential information; and (C) Not use, reproduce, or derive benefit from such Custom Developments or custom models for any other purpose, or for the benefit of any other party, without express written authorization from the Contracting Officer."	
***Description of Issue*: *The scope of "Custom Developments" subject to Government exclusive ownership appears broad enough to capture the Contractor's general methodologies, prompt libraries, and process know-how developed incidentally during performance, giving away rights to Contractor methodology*.** The definition of "Custom Developments" is broad enough to encompass any configuration, enhancement, or workflow modification developed under the contract. Paragraph (e)(4)(viii) dedicates all Custom Developments to the Government's exclusive use and prohibits the Contractor from using, reproducing, or deriving benefit from them for any other purpose without written Contracting Officer consent. This could reach the Contractor's general prompt engineering practices, reusable workflow templates, and internal process improvements that are developed incidentally during contract performance and applied across multiple engagements. Such general	***Proposed Clarification or Recommendation*:** We recommend that GSA revise paragraph (e)(4)(viii) to align the Government's intellectual property rights in Custom Developments with the established FAR/DFARS data rights framework. The Government should generally retain unlimited rights in Custom Developments — including the right to use, modify, reproduce, and redistribute – but the Contractor should retain concurrent rights to use, reproduce, and derive benefit

methodologies and process know-how do not incorporate Government Data or Government confidential information and are not contract-specific deliverables; subjecting them to Government exclusive ownership is inconsistent with established intellectual property procurement law and would discourage Contractor investment in AI innovation.	from such Custom Developments for any purpose without express written Contracting Officer authorization, consistent with the unlimited rights construct under FAR 52.227-14 and DFARS 252.227-7013. GSA should also narrow the definition of "Custom Developments" subject to Government ownership to items specifically identified as deliverables, expressly excluding the Contractor's reusable consulting methodologies, implementation frameworks, prompt libraries developed independently by the Contractor that contain no Government Data, orchestration logic, workflow templates, connectors, evaluation approaches, hosting infrastructure, agents, skills, agentic harnesses and other pre-existing or independently developed intellectual property — even where used in performing Government work.

Comment #8	
***Reference*: *552.239*-7001 Paragraph (e)(4)(ix), Feedback** ***Language at issue*:** "The Government retains ownership of all feedback provided by the Government to the Contractor with respect to the LLM or Custom Developments, regardless of whether such feedback is generated by Government personnel, the Contractor, or through automated processes. If feedback includes Government Data and Government confidential information, that feedback **may not be used** for system improvement purposes or any other purposes (except solely in the performance of this contract)." (emphasis added).	
Description of Issue: ***The feedback ownership and use restriction is fundamentally inconsistent with standard commercial technology licensing terms and would prohibit Contractors from using automated telemetry to remediate system defects in their commercial products.***	***Proposed Clarification or Recommendation***: We recommend that GSA strike paragraph (e)(4)(ix) in its entirety to align with standard commercial

Paragraph (e)(4)(ix) provides that the Government retains ownership of all feedback — including feedback "generated through automated processes" – and prohibits the use of such feedback for system improvement purposes if it includes Government Data or Government confidential information. This provision is fundamentally inconsistent with standard commercial technology licensing terms, which universally permit the licensor to use customer feedback (including thumbs up/down ratings, automated telemetry, and error reporting) to improve the product for the benefit of users. Under the current language, a Contractor that identifies a system defect through automated processes during Government contract performance would be prohibited from using that information to remediate the defect in its commercial product, even where the remediation contains no Government Data. This outcome undermines product safety and quality for all users, including the Government itself.	technology licensing terms or, at a minimum, expressly permit Contractors to use aggregated, de-identified feedback data for general model improvement, unless customers explicitly opt out.

Comment #9	
***Reference*: *552.239*-7001** Paragraph (f)(5), Contractor Obligations for Compliance, Reporting, and Documentation. ***Language at issue*:** "The Contractor must: ... (5) Notify the Contracting Officer and any Government provided point of contacts as soon as possible but not longer than **within 72 hours of the discovery** of **any incident** (**as defined by the Federal Information Security Modernization Act** of 2014 in 44 U.S.C. 3552(b)(2)) affecting any contractors, including Third-party entities, handling Government Data and provide daily status updates to all points of contact until resolved. . . . (ii) Where FedRAMP incident communication and response procedures conflict with contractual requirements, the FedRAMP procedures take priority if the system is required to be FedRAMP Authorized. (iii) The Contractor must preserve all relevant logs, forensic images, and incident artifacts for a minimum of 90 calendar days from a security incident involving Government Data to support follow-on investigation activity by law enforcement entities. (iv) The Contractor must also complete the CISA incident reporting form (*https://myservices.cisa.gov/irf*)." (emphasis added).	
Description of Issue: ***Paragraph (f)(5) imposes incident reporting obligations that are duplicative of existing federal cybersecurity and privacy reporting frameworks and fails to define what constitutes a reportable AI- or LLM-specific incident as distinct from a general cybersecurity incident.*** Paragraph (f)(5) creates duplicative and potentially conflicting incident reporting obligations that overlap with existing federal cyber, privacy, and security reporting frameworks, and fails to define what constitutes a reportable AI- or LLM-specific incident as distinct from a general cybersecurity incident. Federal contractors are already subject to incident reporting obligations under multiple overlapping frameworks, including DFARS 252.204-7012 (cybersecurity incident reporting), FedRAMP Continuous Monitoring requirements, Cyber Security and Infrastructure Security Agency	***Proposed Clarification or Recommendation***: We recommend that GSA harmonize the incident reporting obligations under paragraph (f)(5) with existing Federal reporting frameworks. Specifically, GSA should establish a single, harmonized reporting trigger that coordinates with obligations under new FAR Part 40, DFARS 252.204-7012, FedRAMP Continuous Monitoring, and CISA incident reporting directives, so that

(CISA) incident reporting directives, and agency-specific privacy breach notification procedures (and soon, under forthcoming FAR controlled unclassified information (CUI) incident reporting provisions). Paragraph (f)(5) layers an additional 72-hour notification obligation on top of these existing requirements without harmonizing the reporting triggers, timelines, or points of contact. This creates a risk of duplicative, inconsistent, and potentially conflicting reporting obligations where a single incident triggers notifications under multiple frameworks. Moreover, paragraph (f)(5) defines a reportable "incident" by reference to the Federal Information Security Modernization Act, which addresses general cybersecurity incidents. The clause does not define what constitutes a reportable AI- or LLM-specific incident as distinct from a general cybersecurity event. In the absence of a clear definition, Contractors lack guidance on the scope of reportable events, and Contracting Officers lack a consistent standard for evaluating compliance.	a single incident does not generate duplicative notifications to multiple Government points of contact under separate frameworks with conflicting timelines. GSA should also define reportable AI/LLM-specific incidents as a distinct category from general cybersecurity incidents, with reporting criteria tailored to the risks unique to LLM-enabled systems.

B. Contractor Accountability

Comment #1	
Reference: 552.239-7001 Paragraph (d), Contractor Responsibility for LLM. ***Language at issue:*** "The Contractor must: (1) Exercise due diligence in selecting and overseeing the LLM's Developer(s), System Operator(s), System Integrator(s), and Service Provider(s). (2) Notify the Contracting Officer, within 72 hours, of any known non-adherence to this clause."	
Description of Issue: ***The 72-hour notification requirement for "any known non-adherence" is unreasonably broad, capturing immaterial procedural deviations with no bearing on Government Data security. Additionally, the "overseeing" obligation under paragraph (d)(1) and other sections of the clause, including "transparency and auditability" impose requirements that exceed the Contractor's ability to obtain information from upstream LLM providers.*** Paragraph (d)(2) requires the Contractor to notify the Contracting Officer within 72 hours of "any known non-adherence to this clause." The scope of this notification obligation is unreasonably broad. The clause contains numerous procedural, documentation, and operational requirements, and a literal reading of paragraph (d)(2) would require notification for any deviation (however minor or immaterial) from any provision of the clause, including administrative or procedural requirements with no bearing on the security or integrity of Government Data. This would generate a significant volume of notifications that are neither useful to the Government nor proportionate to the underlying risk, burdening both industry and Contracting Officers. The clause elsewhere contains a 72-hour	***Proposed Clarification or Recommendation***: We recommend that GSA remove the 72-hour notification requirement in paragraph (d)(2) or narrow to instances of non-adherence that involve, or reasonably pose a risk to, the security, integrity, or confidentiality of Government Data. Additionally, GSA should clarify that any oversight obligations apply only to the portions of the LLM system that the Contractor directly develops, configures, or controls. Where the Contractor operates as an LLM User or System Integrator, the Contractor should not be required to

notification requirement for incidents and a separate notice requirement where the LLM has been modified or configured to comply with non-US statutes, regulations, or policies. These and existing Contractor disclosure requirements are sufficient to protect the Government's interests. Separately, paragraph (d)(1) requires the Contractor to "[e]xercise due diligence in selecting and *overseeing*" the LLM's Developer(s), System Operator(s), System Integrator(s), and Service Provider(s). This "obligation" — and related requirements elsewhere in the clause, such as the Transparency and Auditability provision under paragraph (f)(7)(vi) – is impracticable where the Contractor operates as an LLM User or System Integrator. In that capacity, the Contractor does not have access to the underlying model architecture, training data, or internal processing logic of the LLM, as such information is proprietary to the LLM Developer and unavailable to downstream customers. The Contractor should not be held to oversight or compliance obligations requiring it to expose or certify information that is outside its possession or control.	expose, certify, or provide information regarding the underlying model architecture, training data, or internal processing logic of an LLM that is proprietary to the LLM Developer and not made available to downstream customers.

Comment #2	
***Reference*: *552.239*-7001 Paragraph (j)(1)(i), Unbiased AI Principles** ***Language at issue*:** "The Contractor must ensure the LLM is developed and monitored in accordance with the following unbiased AI principles: (i) The LLM **must be truthful** in responding to user prompts seeking factual information or analysis. The LLM must prioritize historical accuracy, scientific inquiry, and objectivity and must acknowledge uncertainty where reliable information is incomplete or contradictory." (emphasis added).	
Description of Issue*: *The absolute "must be truthful" standard under the Unbiased AI Principles is unachievable given the inherent limitations of current LLM technology, including hallucination and factual error, and exposes Contractors to noncompliance liability for inherent technological characteristics. Paragraph (j)(1)(i) requires that the LLM "must be truthful" and "must prioritize historical accuracy, scientific inquiry, and objectivity." This imposes an absolute performance standard. No currently available LLM can guarantee truthful, accurate outputs in all circumstances — hallucination and factual error are inherent, well-documented limitations of the technology. An absolute "must" standard exposes the Contractor to noncompliance liability (including decommissioning costs under paragraph (j)(3)) for behavior that is an inherent characteristic of the technology rather than a failure of Contractor diligence. The standard should be revised to a commercially reasonable efforts obligation.	***Proposed Clarification or Recommendation***: We recommend that GSA revise paragraph (j)(1)(i) to replace the absolute "must" standard with a commercially reasonable efforts obligation (e.g., "The Contractor shall use commercially reasonable efforts to ensure the LLM prioritizes truthfulness, historical accuracy, scientific inquiry, and objectivity").

Comment #3	
***Reference*: *552.239*-7001 Paragraph (j)(1)(ii); (j)(3), Unbiased AI Principles** ***Language at issue*:** "The LLM **must be a neutral, nonpartisan tool** that does not manipulate responses in favor of ideological dogmas. The Contractor must not intentionally introduce or embed partisan or ideological judgments into the LLM's Data Outputs through methods such as training data selection, fine-tuning, Retrieval-Augmented Generation (RAG) references, system prompts, or other configuration methods. . . . (3) Noncompliance. . . . (ii) The Contractor is liable for reasonable decommissioning costs if the Contracting Officer terminates this contract or task/delivery order for cause for failure to remediate after receiving specific written notice of non-compliance from the Contracting Officer with the Unbiased AI Principles described in paragraph (j)(1)." (emphasis added).	
Description of Issue: ***The "neutral, nonpartisan" standard under the Unbiased AI Principles relies on subjective, untestable criteria, and the associated decommissioning liability can attach to base model behavior outside the Contractor's control.*** Paragraph (j)(1)(ii) requires the LLM to be a "neutral, nonpartisan tool" that does not "manipulate responses in favor of ideological dogmas." These terms are inherently subjective and lack objective, testable criteria, creating a risk of inconsistent enforcement across Contracting Officers. Paragraph (j)(3)(ii) compounds this concern by exposing the Contractor to decommissioning cost liability upon termination for failure to remediate noncompliance with the Unbiased AI Principles. The neutrality, accuracy, and absence of ideological bias required by paragraph (j)(1) are substantially a function of the underlying model's training data, model weights and base configuration, not the integrator's system prompts or RAG setup. A System Integrator that has complied with every configuration requirement can nonetheless be exposed to liability if the base model exhibits behavior deemed noncompliant.	***Proposed Clarification or Recommendation***: We recommend that GSA replace the subjective terms "neutral, nonpartisan" and "ideological dogmas" with objective, testable criteria or reference an established evaluation framework. Additionally, because only the LLM Developer can satisfy these requirements at the model level liability should require a demonstrated nexus between the noncompliance and the Contractor's own configuration, action, or omission.

Comment #4
Reference*: *552.239*-7001 Paragraph (j)(3)(ii)(A), Decommissioning liability** ***Language at issue*:** "(3) Noncompliance. . . . (ii) The Contractor is liable for reasonable decommissioning costs if the Contracting Officer terminates this contract or task/delivery order for cause for failure to remediate after receiving specific written notice of non-compliance from the Contracting Officer with the Unbiased AI Principles described in paragraph (j)(1). (A) Decommissioning cost liability are **not to exceed *[Contracting Officer to insert a percentage] at a percentage of contract value." (emphasis added).

Description of Issue: ***The blank decommissioning cost cap, undefined "failure to remediate" standard, and absence of a specified cure period preclude offerors from evaluating risk and invite inconsistent application.*** Paragraph (j)(3)(ii)(A) provides that decommissioning cost liability shall "not exceed __ [*Contracting Officer to insert a percentage*] at a percentage of contract value." Leaving the cost cap as a blank to be completed by the Contracting Officer at time of award precludes offerors from evaluating this risk during the proposal process and invites inconsistent application across contracts. Additionally, the clause does not define "failure to remediate." It is unclear whether this determination is subject to an objective standard or rests solely within the Contracting Officer's discretion. The clause also does not establish a defined cure period within which the Contractor must remediate before decommissioning liability attaches, creating uncertainty as to the timeline for compliance.	***Proposed Clarification or Recommendation***: We recommend that GSA establish a standardized decommissioning cost cap within the clause (or prescribe a ceiling methodology), and define "failure to remediate" by reference to objective criteria including a specified cure period (e.g., 30 or 60 calendar days from receipt of written notice) before decommissioning liability attaches. Decommissioning liability should be limited to instances where the noncompliance is attributable to the Contractor's own conduct, knowledge, or action.

3. Are the roles and responsibilities of the contractor, LLM Developer, LLM System Operator, LLM System Integrator, and LLM Service Provider clearly defined and flowdown paragraphs accurately presented?

Coalition members have raised the following concerns regarding the delineation of Contractor responsibilities relative to those of the LLM Developer, System Operator, and System Integrator:

Comment #1	
Reference: 552.239-7001 Paragraph (d)(3), Contractor Responsibility for LLM ***Language at issue:*** "(3) Demonstrate compliance. The Contractor may satisfy due diligence by relying on: (i) Flowdown of applicable requirements of this clause to LLM's Developer(s), System Operator(s), System Integrator(s), and Service Provider(s); or (ii) Obtaining attestation, from an individual with appropriate authority representing the LLM's Developer(s), System Operator(s), System Integrator(s), and Service Provider(s), that requirements have been implemented."	
Description of Issue: Paragraph (d)(3) provides two mechanisms by which a Contractor may demonstrate due diligence: (i) flowing down applicable clause requirements to the LLM's Developer(s), System Operator(s), System Integrator(s), and Service Provider(s), or (ii) obtaining written attestation from an authorized	***Proposed Clarification or Recommendation***: We recommend that GSA revise paragraph (d)(3) to expressly limit the Contractor's compliance obligations to matters within its reasonable

representative of each such entity that the applicable requirements have been implemented. While the provision of alternative compliance pathways reflects a constructive effort to limit the Contractor's direct compliance burden, neither mechanism is practicable in the commercial AI marketplace. Major commercial LLM providers operate under standardized, non-negotiable terms of service and enterprise agreements that do not accommodate customer-imposed regulatory flowdowns. These providers similarly do not issue individualized compliance attestations to downstream customers on a contract-by-contract basis. As a result, the Contractor is placed in a compliance posture that is dependent on the cooperation of upstream commercial providers over whom it has no contractual leverage, exposing the Contractor to noncompliance risk for circumstances outside its reasonable control.	control. Specifically, GSA should permit the Contractor to satisfy its due diligence obligation by relying on commercially available compliance documentation from upstream providers — including published security certifications (e.g., SOC 2 reports, ISO 27001 certifications), FedRAMP authorization packages, and publicly available terms of service – in lieu of requiring negotiated contractual flowdowns or individualized written attestations. Where the Contractor has made commercially reasonable efforts to obtain the cooperation of the applicable LLM Developer(s), System Operator(s), System Integrator(s), or Service Provider(s) and such cooperation is not provided, the Contractor should not be deemed noncompliant solely by reason of the upstream provider's refusal to accept flowdown terms or issue individualized attestations. Alternatively, we recommend the GSA negotiate Government terms with the LLM/API frontier companies it approves for use and put them on a list of Accepted LLM/API companies for Contractors to consider, eliminating the responsibility for Contractors to do so.

Comment #2
***Reference*: *552.239*-7001 Paragraph (e)(4), Government Data Handling and Processing Requirements.** ***Language at issue*:** "Unless otherwise expressly authorized by the Contracting Officer: . . . (vii) Upon completion, termination or expiration of the contract or task/delivery order, unless otherwise directed in writing by the Contracting Officer, the Contractor must securely and permanently

delete all such Government Data and any Custom Developments from the LLM and all its other systems and all copies, backups and derivatives thereof, and certify deletion to the Contracting Officer in writing..”	
Description of Issue: ***The absolute data deletion requirement under paragraph (e)(4)(vii) is commercially impracticable because it does not account for standard backup retention cycles, disaster recovery architectures, legal hold obligations, or the operational realities of provider-managed cloud infrastructure.*** Paragraph (e)(4)(vii) requires the Contractor, upon completion, termination, or expiration of the contract, to “securely and permanently delete all such Government Data and any Custom Developments from the LLM and all its other systems and all copies, backups and derivatives thereof” and certify deletion in writing to the Contracting Officer. This absolute deletion requirement does not account for several commercially standard and legally required data retention practices. First, enterprise cloud infrastructure and managed service providers maintain automated backup and disaster recovery systems that retain data copies on defined retention cycles; the Contractor cannot unilaterally compel immediate purging of all backup instances from provider-managed infrastructure. Second, the Contractor may be subject to independent legal obligations — including litigation holds, regulatory retention requirements, and audit preservation duties – that prohibit deletion of certain data during applicable retention periods. Third, the requirement to delete “all copies, backups and derivatives thereof” is operationally impracticable in distributed cloud environments where data may be replicated across multiple availability zones and backup tiers as part of standard infrastructure resilience.	***Proposed Clarification or Recommendation:*** We recommend that GSA revise paragraph (e)(4)(vii) to (1) require the Contractor to use commercially reasonable efforts to delete Government Data and Custom Developments, recognizing that it is impracticable for the Contractor to identify and purge Government Data from all systems where such data may reside, including systems maintained by third-party providers over which the Contractor does not exercise direct control; (2) expressly exempt from the immediate deletion requirement any copies of Government Data retained in automated backup and disaster recovery systems, provided such copies are deleted in accordance with the Contractor’s standard backup retention cycle and are not affirmatively accessed or used for any purpose following contract completion; and (3) permit the Contractor to retain Government Data to the extent required by applicable law, regulation, or legal hold obligations, provided the Contractor notifies the Contracting Officer of the retention, identifies the legal basis, and continues to protect such data in accordance with the clause’s data handling requirements.

Comment #3

***Reference*: 552.239**-7001 Paragraph (f)(1), Contractor Obligations for Compliance, Reporting, and Documentation

***Language at issue*:** "The Contractor must: . . . (1) Disclose all LLMs used **or made available** in performance of this contract or a task/delivery order. Disclose all entities filling the roles defined by the flowdown supplemental clauses. Disclosures are due to the Contracting Officer or ordering Contracting Officer by the date specified in the contract or task/delivery order. If no date is specified, disclosures are due within 120 days after commencing work under the contract or task/delivery order." (emphasis added).

Description of Issue:	***Proposed Clarification or Recommendation:***
The LLM disclosure requirement under paragraph (f)(1) is overbroad because the phrase "or made available" extends the Contractor's reporting obligation beyond LLMs actually used in contract performance to encompass every LLM in the Contractor's enterprise environment, creating an administratively unworkable obligation for both Contractors and the Government. Paragraph (f)(1) requires the Contractor to disclose "all LLMs used or made available in performance of this contract or a task/delivery order." The phrase "or made available" is vague and materially expands the disclosure obligation beyond LLMs actually used in contract performance. Read literally, this language would require the Contractor to identify and report every LLM to which its personnel have access within the enterprise environment — regardless of whether the LLM was used in connection with the Government contract. For large enterprises with multiple commercial AI tools deployed across business units, this would generate voluminous disclosures encompassing LLMs with no nexus to Government work. Such disclosures are administratively burdensome for both the Contractor and the Government: the Contractor must compile and maintain an inventory of all available LLMs irrespective of their relevance to the contract, and the Contracting Officer must review disclosures that do not meaningfully inform the Government's understanding of which LLMs are actually being applied to federal engagements. The disclosure obligation should be limited to LLMs used in performance of the contract, which is sufficient to achieve the Government's transparency objectives without imposing disproportionate administrative burden on either party.	We recommend that GSA remove the phrase "or made available" from paragraph (f)(1) and limit the Contractor's disclosure obligation to LLMs actually used in performance of the contract or task/delivery order.

Comment #4

***Reference*:** 552.239-7001 Paragraph (f)(7)(iii), (viii), and (xii), Contractor Obligations for Compliance, Reporting, and Documentation

***Language at issue*:** "The Contractor must: . . . (7) Provide, upon Government request, and under appropriate confidentiality protections, existing commercial documentation or disclosures sufficient to demonstrate compliance, including but not limited to: (iii) LLM decision-making

processes, logic, and operational parameters; . . . (viii) Testing methodologies used to detect and mitigate noncompliance with the unbiased AI principles described in paragraph (j)(1); (ix) Known biases (including commercial, political, or personal considerations, advertising, endorsements, or fraudulent/corrupt interests), limitations, truthfulness concerns, and performance metrics. . . and (xii) Any other information necessary for the Government to monitor and evaluate the LLM's performance, risks, and effectiveness and to complete an LLM Impact Assessment required by OMB M-25-21 and successor requirements." (emphasis added).

Description of Issue: ***The compliance documentation requirements under paragraph (f)(7) require the Contractor to produce information about the LLM's internal architecture, testing methodologies, and performance characteristics that is proprietary to the LLM Developer and unavailable to Contractors operating as LLM Users or System Integrators.*** Paragraph (f)(7) requires the Contractor, upon Government request, to provide documentation sufficient to demonstrate compliance with the clause. Several categories of required documentation — including the LLM's "decision-making processes, logic, and operational parameters" (paragraph (f)(7)(iii)), "testing methodologies used to detect and mitigate noncompliance with the unbiased AI principles" (paragraph (f)(7)(viii)), and information necessary to "complete an LLM Impact Assessment required by OMB M-25-21" (paragraph (f)(7)(xii)) – call for information that resides exclusively with the LLM Developer and is not disclosed to downstream customers. A contractor operating as an LLM User or System Integrator has no access to the underlying model's architecture, weighting logic, training data provenance, or internal testing methodologies, and cannot produce documentation it does not possess. The Contractor can document its own implementation — including prompt configuration, RAG pipeline validation, and output sampling – but should not bear compliance risk for failing to furnish proprietary information that the LLM Developer does not make available. Additionally, paragraph (f)(7)(xii) references an "LLM Impact Assessment," whereas OMB M-25-21 defines an "AI Impact Assessment." This terminology should be corrected or the reference removed, and the Contractor's obligation to support such an assessment should be limited to information within its possession or control.	***Proposed Clarification or Recommendation:*** We recommend that GSA revise paragraph (f)(7) to (1) limit the Contractor's documentation, transparency, and auditability obligations to information within the Contractor's possession or control, and expressly permit the Contractor to satisfy such obligations by furnishing commercially available documentation from upstream LLM providers; and (2) correct the reference to OMB M-25-21 in paragraph (f)(7)(xii) for the term "LLM Impact Assessment."

Comment #5

Reference: 552.239-7001 Paragraph (h), Data Portability and Interoperability.

Language at issue: "The Contractor must: (1) Provide and ensure that LLM services support the export of Government Data (including but not limited to: user-generated content, interaction history, uploaded content and media, and knowledge bases) in open or standard machine-readable formats, together with **sufficient associated metadata and schema** information necessary for the Government to retrieve, use, or migrate Government Data to another service. (2) Provide and support documented APIs sufficient for Government integration and data export to the extent

the service uses APIs. (3) Not impose proprietary formats, technical restrictions, **additional costs**, or additional licensing conditions that materially impair the Government's ability to retrieve, use, or migrate Government Data to another service. (4) Provide tooling or interfaces to facilitate migration of Government Data and Custom Developments in a secure manner."	
Description of Issue: ***The data portability and interoperability requirements under paragraph (h) are overbroad because they effectively compel architectural unbundling, non-standard data export formats, and disclosure of proprietary schemas and business logic for enterprise platforms that incorporate LLM functionality as one component of a broader service offering.*** Paragraph (h) requires the Contractor to support export of Government Data in open or standard machine-readable formats, provide documented APIs for Government integration and data export, refrain from imposing proprietary formats or technical restrictions that impair the Government's ability to migrate Government Data, and provide tooling or interfaces to facilitate migration of Government Data and Custom Developments. While the Government's interest in data portability is reasonable, these requirements do not account for the architectural realities of modern enterprise software platforms. Enterprise platforms frequently incorporate LLM functionality as one layer within a broader service offering that includes proprietary platform services, third-party components, customer-configurable workflows, and integrated business logic. A literal reading of paragraph (h) would require the Contractor to unbundle its platform architecture, develop non-standard export mechanisms, and disclose proprietary schemas, metadata structures, and business logic solely because LLM functionality is present within the platform — even where the Government's data is otherwise fully exportable through the platform's commercially available export tools. Paragraph (h)(1)'s requirement to provide "sufficient associated metadata and schema information" is particularly problematic, as it could be read to compel disclosure of proprietary data models and internal system architectures that constitute the Contractor's trade secrets and competitive intellectual property. The data portability obligation should be limited to Government Data itself and should permit the Contractor to satisfy its export obligations through commercially available export mechanisms rather than requiring personalized tooling or architectural modifications. Finally, the prohibition on "additional costs" under Paragraph (h)(3) is commercially unreasonable insofar as it would preclude the Contractor from recovering legitimate costs associated with data export, migration support, or the development of export functionality beyond the Contractor's standard commercial offering. The Contractor should not be required to absorb the cost of non-standard export or migration services at no charge to the Government.	***Proposed Clarification or Recommendation:*** We recommend that GSA revise paragraph (h) to align the Contractor's data portability obligations with commercially reasonable export capabilities. Specifically, GSA should: (1) clarify that the Contractor may satisfy its data export, format, and migration tooling obligations under paragraph (h) through commercially available export tools, standard platform functionality, and standard commercial data formats, and is not required to develop tailored export mechanisms, custom migration tooling, or unbundle proprietary platform architectures solely because LLM functionality is incorporated within the platform; (2) limit the scope of the "metadata and schema information" requirement under paragraph (h)(1) to metadata necessary to render the exported Government Data usable by the Government, and clarify that it does not require disclosure of the Contractor's proprietary internal system architecture, data models, or business logic; and (3) and (3) remove "additional costs" from the list of prohibited impositions under paragraph (h)(3).

Comment #6	
Reference: 552.239-7001 Paragraph (i)(1) and (i)(3), Change Notification ***Language at issue:*** "(1) Advance Notice of Material Changes. The Contractor must **provide written notice to the Contracting Officer at least thirty (30) calendar days prior to any planned material change affecting the services** provided under this contract, including: (i) Adding, replacing, or materially changing any LLMs, LLM's Developer(s), System Operator(s), System Integrator(s), or Service Provider(s); (ii) Changes to the services handling Government Data, including discontinuation or material reduction of data protection controls applied to Government Data, or change in FedRAMP Authorization Status; or (iii) Any modification, configuration, or training change to the LLM made to comply with any non-U.S. government statute, regulation, or policy. . . . (3) Safety and Performance Degradation. The Contractor must notify the Contracting Officer **within seven (7) calendar days of identifying any change** that materially increases output bias; decreases safety guardrails or behavioral constraints; or degrades performance or truthfulness of outputs. The notification must describe the change, its purpose, the evaluation approach used and any new limitations, trade-offs, or potential negative impacts identified.. ." (emphasis added).	
Description of Issue: ***Obligations surrounding changes to the LLM or Developer, Operator, Integrator, or Service Provider, and notice and safety degradation are impracticable for prime Contractors using third-party foundation models over which they have no visibility or control.*** Paragraphs (i)(1) and (i)(3) impose change notice obligations that are impracticable for a prime Contractor using a third-party foundation model. Paragraph (i)(1) requires 30 calendar days' advance written notice before any material change to the LLM or its Developer, Operator, Integrator, or Service Provider. Similarly, Paragraph (i)(3) requires notification within 7 calendar days of identifying any change that materially increases output bias, decreases safety guardrails, or degrades performance. A prime Contractor using a commercial foundation model has no contractual leverage to compel advance notice from the model developer and will frequently have no knowledge of model changes unless the vendor publishes release notes. Paragraph (i)(2) already applies a "reasonable efforts" standard to model version changes, but paragraphs (i)(1) and (i)(3) impose absolute obligations. The prime Contractor should not bear liability for failing to provide notice of changes about which it has no knowledge and no practical ability to discover.	***Proposed Clarification or Recommendation***: GSA should align obligations with the entity's actual role and control over the LLM system, data processing environment, and security controls. We recommend that GSA extend the "reasonable efforts" standard currently applied in paragraph (i)(2) to the notice obligations in paragraphs (i)(1) and (i)(3). Both provisions should be qualified to require notice "to the extent such information is known to, or reasonably discoverable by, the Contractor based on information disclosed by the applicable LLM Developer or System Operator."

Comment #7
Reference: 552.239-7001 Paragraph (j)(1)(iii), Unbiased AI principles.

Language at issue: "(iii) The Contractor must implement continuous improvement processes to enhance detection and mitigation of performance, trustworthiness, bias, and/or systems generating illegal or prohibited content, including regular evaluation of system outputs (excluding Data Outputs) against verified factual sources."	
Description of Issue: ***The continuous improvement and verified factual source evaluation requirements under paragraph (j)(1)(iii) are impracticable for Contractors operating as LLM Users, System Operators, or System Integrators because they require modification of base-model behavior and training processes that are exclusively within the LLM Developer's control.*** Paragraph (j)(1)(iii) requires the Contractor to "implement continuous improvement processes to enhance detection and mitigation of performance, trustworthiness, bias, and/or systems generating illegal or prohibited content, including regular evaluation of system outputs (excluding Data Outputs) against verified factual sources." This obligation presupposes that the Contractor has the ability to modify the LLM's underlying model behavior — including its factual accuracy, bias characteristics, and content generation parameters – through changes to training data, model weights, or internal processing logic. However, contractors operating as LLM Users, System Operators, or System Integrators typically do not train, own, or control the base-model weights of third-party foundation models. In that capacity, the Contractor cannot modify a hosted third-party base model to alter its factual accuracy, bias profile, or content generation behavior, nor can the Contractor conduct fact-verification processes that change the model's underlying outputs. These capabilities reside exclusively with the LLM Developer that trains and maintains the model. A literal reading of paragraph (j)(1)(iii) would require the Contractor to build and maintain a Government-specific variant of each commercial model that diverges from the commercially available service — an outcome that is inconsistent with the Administration's policy of leveraging commercial-item solutions and that would impose disproportionate cost and complexity on both industry and the Government.	***Proposed Clarification or Recommendation***: We recommend that GSA revise paragraph (j)(1)(iii) to scope the Contractor's continuous improvement obligations to measures within the Contractor's possession or control. Specifically, where the Contractor operates as an LLM User, System Operator, or System Integrator, the Contractor's obligation should be limited to implementing continuous improvement processes at the application layer rather than requiring modification of the base model's training data, model weights, or internal processing logic. Additionally, the requirement for "regular evaluation of system outputs (excluding Data Outputs) against verified factual sources" should be qualified as a commercially reasonable efforts obligation, recognizing that comprehensive fact-verification of all system outputs is not operationally feasible at scale.

4. Do you understand how to implement the flowdown clauses?

While the flowdown requirements are generally clear, Coalition members continue to have the following concern:

Comment #1
Reference: 552.239-7001 Paragraph (a)(2), Flowdown Requirements

Language at issue: "The Contractor must extend (flowdown) the specific paragraphs of the base clause to any subcontractor or service provider (referred to as "Contractor" in each flowdown supplemental clause) functioning in the applicable roles, as defined in the flowdown supplemental clauses. Where a single entity performs multiple roles, multiple flowdown supplemental clauses should be used: (i) 552.239-7001-1 LLM Developer Flowdown Requirements (DATE); (ii) 552.239-7001-2 LLM System Operator Flowdown Requirements (DATE); (iii) 552.239-7001-3 LLM System Integrator Flowdown Requirements (DATE); (iv) 552.239-7001-4 LLM Service Provider Flowdown Requirements (DATE)."	
Description of Issue: ***Flowdown requirements are impracticable for enterprise-licensed commercial tools.*** The flowdown requirements create an impracticable compliance obligation with respect to enterprise-licensed commercial tools that were not procured specifically for contract performance. Where the clause appropriately applies to a contract based on the scope of work, and the Contractor utilizes its general-purpose enterprise AI tools in the course of performance, those commercial tool providers could technically qualify as LLM Service Providers or LLM Developers subject to flowdown. However, the Contractor has no realistic mechanism to impose the clause's flowdown terms on such vendors, whose commercial agreements are non-negotiable. This creates an impossible compliance position for the Contractor.	***Proposed Clarification or Recommendation***: GSA should exclude from the flowdown requirements any commercial tool vendor whose product constitutes the Contractor's general-purpose enterprise capability and was not procured specifically for performance of the contract.

5. Does the clause adequately address risks related to foreign ownership or control of LLMs, where changes to the LLM could covertly affect Government Data, outputs, or decisions without changing the contracting entity?

Coalition members have raised the following concerns related to the scope of foreign contributions to the LLM:

Comment #1	
Reference: 552.239-7001 Paragraph (f)(2)(iii), Contractor Obligations for Compliance, Reporting, and Documentation - Component Flexibility & Risk Mitigation ***Language at issue:*** "**Incidental foreign-developed components** (*e.g.,* open-source components, published research), ancillary Third-party services, or globally operated infrastructure dependencies are permissible, provided: (A) They do not introduce security risks or foreign control that would violate criteria (f)(2)(i) or (f)(2)(ii). (B) The systems storing or processing Government Data satisfy applicable Federal security requirements, including those for data residency and isolation. (C) A risk-based approach is applied, focusing on objective criteria such as ownership, control, hosting, and security posture." (emphasis added).	
Description of Issue: ***The scope of incidental foreign-developed components is unclear.***	***Proposed Clarification or Recommendation***: We request that GSA

Paragraph (f)(2)(iii) permits "incidental foreign-developed components (e.g., open-source components, published research)" provided they do not introduce security risks or foreign control. However, it is unclear whether this includes open-weight or open-source foundation models with foreign contributors. Many widely used open-source LLMs involve contributions from developers in multiple countries, and it is unclear whether foreign code contributions to an otherwise domestically maintained open-source model would trigger disqualification or require additional risk assessment under paragraphs (f)(2)(i) or (f)(2)(ii). The absence of guidance on this point creates uncertainty for contractors evaluating whether a given open-source model is permissible for use on Government contracts.	provide guidance on the treatment of open-weight and open-source LLMs with foreign contributors under paragraph (f)(2)(iii). Specifically, does an incidental foreign code contribution to an open-source model—where the model is maintained by a U.S.-based organization and the foreign contribution does not confer foreign ownership or control–trigger disqualification under paragraphs (f)(2)(i) or (f)(2)(ii)? We recommend that GSA clarify that the risk-based assessment under (f)(2)(iii)(C) applies to open-source models and that incidental foreign contributions do not, standing alone, constitute foreign control or introduce a *per se* security risk.

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
Coalition for Government Procurement