

April 3, 2026

Multiple Award Schedule Program Management Office
Federal Acquisition Service
U.S. General Services Administration
1800 F Street NW
Washington, D.C., 20006

Re: Comments on Proposed GSAR Clause 552.239-7001, Basic Safeguarding of Artificial Intelligence Systems, Included in MAS Refresh 31

To Whom It May Concern:

TechNet appreciates the opportunity to provide comments on the General Services Administration's (GSA) proposed addition of General Services Acquisition Regulation (GSAR) clause 552.239-7001, *Basic Safeguarding of Artificial Intelligence Systems*, as part of the Multiple Award Schedule Refresh 31. TechNet represents a broad network of leading technology companies committed to advancing responsible AI development and deployment across the economy, including in support of federal missions. We strongly support clear and workable federal acquisition frameworks that enable federal agencies to adopt innovative commercial AI solutions while maintaining strong safeguards for security, privacy, and government data.

President Trump and his administration have made accelerating federal adoption of AI a clear priority by directing agencies to rely more heavily on commercially available solutions, reduce unnecessary procurement barriers, and modernize acquisition practices so the government can benefit more quickly from private-sector innovation. That policy direction has created important momentum toward broader federal AI adoption. To continue that momentum, acquisition terms must reflect how commercial AI systems are designed, licensed, updated, and secured in practice. However, several elements of the proposed GSAR clause risk moving in the opposite direction by introducing government-unique requirements that could slow procurement, narrow provider participation, and limit agency access to rapidly evolving AI capabilities. The proposed GSAR clause would require contractors to build and maintain a parallel, government-only product distinct from their commercial products, effectively converting what is nominally a commercial acquisition into a de facto bespoke procurement. The clause is particularly likely to deter nascent and emerging AI providers, who are least able to deviate from standard commercial terms and most likely to be sources of cutting-edge innovation—precisely the companies GSA should be positioned to access.

Alignment with Commercial Acquisition Principles and Existing Security Frameworks

TechNet is concerned that the proposed GSAR clause moves federal AI procurement away from established commercial acquisition principles and toward a bespoke procurement framework that is inconsistent with Federal Acquisition Regulation (FAR) Part 12 and the Administration's broader commercial acquisition posture. The Administration has repeatedly emphasized that federal procurement should default to commercially available solutions wherever possible and avoid imposing government-unique requirements that unnecessarily raise barriers to entry for innovative providers. That principle is reflected in the Revolutionary FAR Overhaul, which is intended to operationalize the statutory commercial item mandate at the regulatory level by reducing unnecessary clauses, narrowing non-commercial requirements, and modernizing acquisition rules so federal agencies can access private-sector innovation more quickly and efficiently. The central purpose of that effort is to ensure agencies are able to buy commercial technology, using market-tested products and services under terms that preserve the economic and operational realities of commercial delivery rather than forcing providers into bespoke federal arrangements.

The MAS mechanism itself is a commercial item contract vehicle, awarded under FAR Part 12. Incorporating a clause that explicitly overrides commercial terms, as paragraph (b) does, is incompatible with the foundational premise of the MAS schedule as a commercial contracting vehicle.

By contrast, the proposed clause would require contractors and service providers to materially restructure commercial AI offerings, alter upstream licensing arrangements, and assume obligations that are neither commercially standard nor operationally scalable. Rather than leveraging commercial market practices, the clause imposes a distinct federal operating model for AI services that many providers cannot implement without fundamentally separating government offerings from their broader commercial platforms. This risks undermining the very objective of current acquisition modernization efforts by reintroducing the type of government-unique contractual burdens that procurement reform is intended to remove.

Commercial AI services are generally offered through standardized terms (like the MAS mechanism), shared infrastructure, continuously updated software environments, and security controls that apply consistently across a broad customer base. Providers maintain these systems at scale, deploying security improvements, performance refinements, and service updates across customer environments in ways that preserve consistency, reliability, and cost efficiency. Requiring vendors to redesign those offerings solely for federal use risks discouraging participation by the very providers best positioned to support secure and advanced government adoption.

This departure from commercial practice is particularly significant because the federal government has already established mature pathways for evaluating cloud-based and AI-enabled services through the Federal Risk and Authorization Management Program (FedRAMP) and related Security Requirements Guide processes. FedRAMP was

specifically designed to create a uniform government-wide security assessment model that allows agencies to rely on a common baseline of controls, documentation, incident response obligations, continuous monitoring, and third-party assessment rather than requiring separate acquisition vehicles to impose distinct security structures. Several provisions in the proposed clause duplicate or expand upon these established frameworks by layering additional incident reporting, log preservation, documentation, audit access, and data handling obligations without clarifying how they interact with existing authorizations. Incident reporting requirements should apply only to clearly defined and material security events, with a threshold high enough to avoid triggering reporting obligations for routine operational activity, preliminary indicators, or low-risk anomalies that do not present meaningful impact to government data or systems. The proposed 72-hour reporting timeline in paragraph (e)(4) is a workable standard that aligns more closely with established defense contracting practice, including DFARS 252.204-7012, and should be preserved rather than shortened through cross-reference or implementation guidance to mirror more compressed reporting timelines used in other federal security frameworks. Finally, the 'eyes off' requirement as currently drafted could prevent contractors from conducting necessary security operations, debugging, and troubleshooting that require human review of logs or data patterns.

This is especially problematic in areas governed by shared responsibility between cloud providers, agencies, and integrators. For example, requiring scheduled contractors to guarantee compliance by upstream service providers, disclose all AI systems used in contract performance, and satisfy new operational transparency requirements may be impractical where commercial providers rely on layered service architectures that already operate within authorized cloud environments. Rather than building on existing security frameworks, the proposed GSAR clause creates a parallel compliance structure that could fragment procurement expectations across agencies and reduce the value of existing authorizations.

Federal AI acquisition policy should align with the spirit of FAR Part 12 and build on the progress already achieved through FedRAMP and SRG processes, not create overlapping requirements that effectively convert commercial cloud and AI services into traditional systems integration procurements that are objectively non-commercial. Preserving alignment with these established frameworks will better support both security and competition while ensuring agencies can continue to access modern commercial AI capabilities at speed.

Intellectual Property and Commercial Licensing

The proposed intellectual property provisions would benefit from further refinement to ensure they reflect how commercial AI systems are developed, licensed, and continuously improved in practice. Several of the core definitions in the GSAR clause are written broadly enough to create uncertainty about ownership of technologies, service data, and derivative improvements that remain integral to privately developed commercial offerings and risk transferring ownership of proprietary IP.

The definition of “Custom Development” is particularly concerning because its inclusion of modifications, customizations, configurations, and enhancements could be interpreted to grant the government ownership interests in fine-tuned models, adapted system behavior, or other model-level changes that arise through use of commercial AI services in a federal environment. Additional clarification is also needed regarding the definition of “Data Inputs,” which as drafted appears broad enough to encompass system prompts and other provider-developed instruction layers that are core to how commercial AI systems function. Because these prompt structures often represent highly sensitive proprietary assets that shape model behavior, safety controls, and performance, they should not be treated as government-owned inputs simply because they are used in connection with contract performance.

In modern AI systems, these forms of refinement are often inseparable from provider-owned model architecture, orchestration layers, and proprietary technical infrastructure. Treating such activities as government-owned work product risks capturing derivative improvements that remain fundamentally dependent on privately developed intellectual property and reflects a broader misunderstanding of how commercial AI services are deployed and maintained. This provision would prevent contractors and service providers from using any generative AI content and solutions developed for a government customer for other engagements even where such work reflects general technical expertise or methodology that does not incorporate government confidential information. This outcome is inconsistent with FAR 52.227-14's careful framework for distinguishing government-funded from privately-funded development.

The treatment of “Data Outputs” and “Government Data” creates similar concern because, when read together, those definitions may extend well beyond agency prompts and generated responses to include metadata, logs, telemetry, audit records, service usage information, and other technical data generated through operation of the service. These categories of information are essential to maintaining reliability, monitoring security, improving performance, and operating AI systems at scale across customer environments. If such information is treated as government-owned data, providers may be restricted from using routine operational data necessary to improve products, maintain service quality, or develop future model capabilities. By prohibiting contractors from leveraging their own innovations across multiple customers, the provision effectively requires custom, government-specific AI development rather than allowing use of commercial solutions. This could create a significant barrier to commercial AI adoption, as contractors would need to maintain entirely separate systems and methodologies for government work, dramatically increasing costs. A narrower definition that clearly excludes service metadata, telemetry, audit logging, and provider business information would better protect agency-controlled content while preserving commercially necessary system operations.

The clause also creates ambiguity around ownership of feedback and service improvements. Commercial AI services evolve continuously through feedback, testing, and system-wide product refinement across customer environments. Provisions suggesting that feedback, enhancements, or derivative developments connected to

contract performance become government-owned create uncertainty around whether providers may continue using ordinary product feedback to improve shared cloud services. Without clarification, providers may be forced to limit feedback mechanisms for federal customers or establish separate operating processes that reduce the government's ability to benefit from ongoing commercial service improvement.

The proposed license framework likewise departs from established commercial practice. Requiring contractors to provide a fully paid-up, royalty-free license does not align with how advanced AI and cloud services are offered in the market, where pricing is generally tied to consumption, compute, storage, or service access over time. For many commercial AI offerings, there is no fully paid-up licensing construct because the service itself depends on continuously provisioned infrastructure and metered usage. In addition, the requirement that providers grant rights allowing the government to copy and store covered systems creates uncertainty regarding whether federal agencies may seek rights extending beyond ordinary service access into underlying owned infrastructure. This license scope is broader than even the Defense Federal Acquisition Regulation Supplement (DFARS) Government Purpose Rights under 252.227-7013, which apply only to data developed with mixed funding, as it could authorize the government to integrate a contractor's proprietary AI system into unrelated government programs, share access with other agencies or contractors, or leverage the system in ways that directly compete with the contractor's commercial offerings.

More broadly, the clause risks displacing commercially negotiated license boundaries that govern scale, integration, access conditions, and usage limits across shared commercial environments. These terms are not incidental to commercial AI services. They are core to how providers manage service integrity, security, pricing, and infrastructure allocation. Final requirements should preserve commercially standard licensing structures and avoid creating a government-specific license model that many providers cannot operationalize within existing service offerings.

Order of Precedence and Flowdown of Commercial Terms

Additional clarity is needed regarding the GSAR clause's order-of-precedence provision and its interaction with existing commercial terms. As drafted, the clause appears to override standard commercial AI terms and other negotiated licensing conditions, creating uncertainty for providers operating under established commercial agreements and for contractors that rely on upstream service providers whose terms cannot be unilaterally altered. Many Schedule holders, systems integrators, and managed service providers procure AI capabilities through layered commercial arrangements with model developers, cloud providers, and other technology vendors whose licensing structures are standardized across customer environments and are not subject to downstream modification by resellers or integrators. Imposing a clause that supersedes those existing terms risks creating obligations that contractors cannot realistically flow through the commercial supply chain and may ultimately discourage participation by providers that depend on upstream commercial platforms to serve federal customers.

More broadly, this approach conflicts with longstanding federal efforts to reduce barriers for commercial companies seeking to support government requirements through existing market offerings rather than government-unique contractual structures. For years, federal acquisition policy has sought to limit unnecessary deviations from commercial practice so agencies can benefit from broader participation by innovative providers. The breadth of the order-of-precedence provision may also create tension with recent statutory direction, including Section 1824 of the *National Defense Authorization Act for Fiscal Year 2026*, titled *Limitation on Required Flowdown of Contract Clauses to Subcontractors Providing Commercial Products or Commercial Services*, which reflects congressional intent to limit the expansion of government-specific clause flowdown obligations in commercial supply chains, including for Department of Defense orders placed through GSA contracting vehicles. Final requirements should be carefully tailored so they do not unintentionally reintroduce the very commercial contracting burdens federal acquisition policy has sought to reduce.

Compliance, Documentation, and Evaluation Requirements

Several compliance provisions impose obligations that are difficult to operationalize within commercial AI delivery models and in some cases require disclosures that providers cannot reasonably make without exposing highly sensitive technical information. The requirement to provide documentation regarding AI system decision-making processes, logic, and operational parameters goes beyond commercially standard disclosures and risks requiring release of proprietary orchestration methods, system architecture, and internal design choices central to provider intellectual property. To the extent that prime contractors do not have existing commercial documents that can fulfill this requirement, it imposes non-commercial disclosure obligations.

Relatedly, requirements to disclose testing methodologies used to assess compliance with performance and neutrality-related provisions may compel release of internal evaluation methods, benchmarking approaches, or governance structures developed as part of proprietary model assurance programs. These are not standardized commercial disclosures and, absent clearer limits and confidentiality protections, could expose highly sensitive technical practices that providers rely upon to improve system performance and reliability.

The clause also introduces new tooling obligations that extend beyond established federal security and privacy requirements without clear implementation boundaries or recognition of the associated operational cost. Requirements to provide tools enabling the government to maintain detailed records of processing activities involving government data and to implement government-configurable controls to detect, manage, prevent, and reject the entry or persistence of personally identifiable information could require new interfaces, reporting capabilities, and product modifications that are not reflected in current contract pricing and go beyond existing FedRAMP control requirements. As drafted, these provisions are broad enough to create uncertainty regarding whether contractors would be expected to expose

sensitive operational logs, telemetry, or internal security records that are not ordinarily made available through commercial service offerings.

Many providers already support privacy and security protections through existing controls, but translating those capabilities into bespoke government-configurable tooling would impose additional development, maintenance, and onboarding costs that are difficult to operationalize at scale. Where retained, these requirements should be more clearly scoped and aligned with existing federal security frameworks as noted above to avoid creating duplicative obligations that extend beyond current commercial and FedRAMP-based practices.

The requirement that contractors disclose all AI systems used in performance of the contract is also overly broad and will extend to subcontractors, independent software vendors, managed service providers, systems integrators, and upstream service providers whose tools support layered delivery environments. Prime contractors frequently do not possess complete visibility into proprietary technical systems used by all subcontractors, particularly where those tools are used internally rather than directly delivered to the government. Contractors cannot compel third-party model providers to expose proprietary algorithmic processes, restructure commercial products, or maintain government-specific versions of their systems. This provision is fundamentally inconsistent with FAR 52.212-4(c)(2), which gives prime contractors discretion to flow down only a minimal number of additional clauses necessary to satisfy contractual obligations. A narrower standard focused on AI systems materially used to deliver contract requirements would be more practical and better aligned with commercial realities.

The performance and evaluation provisions also incorporate subjective standards that are difficult to define and apply consistently across general-purpose AI systems, while linking those standards to meaningful contractual consequences. Requirements tied to truthfulness, historical accuracy, scientific objectivity, and the absence of ideological judgment are inherently difficult to operationalize in a uniform way across broad model deployments, particularly where outputs may depend on context, source material, prompt structure, or evolving factual interpretation. Concepts such as neutrality, objectivity, and bias are not self-defining technical standards and may vary depending on the evaluator, benchmark design, and use case. As drafted, these provisions risk creating enforceable obligations around concepts that remain inherently subjective and difficult to measure in a predictable procurement context.

This concern is amplified by the clause's evaluation structure, which permits the government to assess systems using agency-developed benchmarks and methodologies while not requiring disclosure of the underlying testing frameworks used to determine compliance. Providers may therefore face remediation demands or performance disputes without a clear understanding of the standards against which systems are being evaluated. In addition, requiring contractors to disclose internal testing methodologies used to assess neutrality or performance would compel release of proprietary model evaluation approaches that many providers treat as highly sensitive intellectual property.

The clause also introduces unnecessary liability exposure through its reference to contractor responsibility for decommissioning costs tied to non-compliance. Existing federal acquisition rules already establish clear remedies when performance issues arise, including termination for cause under FAR Part 12 and default-related remedies under FAR Part 49 where applicable. Adding a separate decommissioning cost obligation on top of those existing authorities creates duplicative and potentially open-ended liability that is not clearly bound by traditional procurement principles. This is particularly problematic in the AI context, where performance findings may rely on evolving benchmarks or subjective standards that are not fully transparent to contractors. Providers should not face additional financial exposure tied to undefined remediation expectations beyond the remedies already contemplated under existing acquisition law. If retained, any transition obligations should be narrowly defined, commercially reasonable, and limited to circumstances involving clear and material failure to meet explicit contractual requirements.

Finally, the GSAR clause creates open-ended exposure by requiring contractors to implement future Office of Management and Budget directives upon request, even where those directives may introduce technical obligations, restrictions, or liabilities that are not defined at the time of contract award. Unlike requirements that can be priced, assessed, and operationalized at the outset of performance, this approach creates uncertainty around future compliance expectations and effectively asks contractors to accept unknown obligations that may materially alter service delivery, product design, or contractual risk over the life of the agreement. A more workable approach would limit this requirement to clearly defined obligations adopted through established procurement processes rather than incorporating future directives that cannot yet be evaluated commercially or operationally.

Change Management and Operational Flexibility

The proposed GSAR clause's change management provisions remain misaligned with the continuous update cycles that define commercial AI services. Requiring thirty days' advance notice before adding or materially changing a service provider, coupled with obligations to provide successor models and notify agencies of changes affecting safety guardrails or output behavior, imposes a static operating model on technologies that improve continuously through rapid iteration. Note that "material" is not defined in the clause, creating uncertainty about which changes trigger notification requirements.

Commercial AI systems are routinely updated through security patches, infrastructure changes, model routing improvements, performance tuning, and service refinements that occur on compressed timelines and across globally shared environments. Imposing broad advance notice obligations for these changes creates administrative burdens that are not scalable and may slow deployment of security or reliability improvements for government users. The clause also creates open-ended exposure by requiring contractors to implement future Office of Management and Budget (OMB) directives upon request, even where those directives may introduce undefined future technical obligations not known at contract award.

Scope, Intended Use, and Applicability of Covered AI Systems

Several provisions in the proposed clause would benefit from additional clarification to ensure that contractual obligations are applied in a manner consistent with how commercial AI systems are designed, deployed, and procured for specific operational purposes. In particular, the requirement that an AI system must not refuse to produce lawful outputs warrants further refinement so that it does not unintentionally require providers to make systems available for uses that extend beyond their intended design, validated deployment environment, or commercial purpose. General-purpose AI systems and purpose-built AI tools operate differently in practice, and procurement requirements should recognize that distinction. Many AI offerings are developed for narrowly defined use cases and are intentionally constrained to support reliability, safety, and performance within those contexts. A system designed to support a specific administrative or operational task, for example, should not be expected to generate outputs unrelated to that function simply because such outputs may be lawful. A prior authorization tool designed to assist with claims review should not be expected to perform unrelated medical inquiry, broad diagnostic support, or other off-domain functions that fall outside its intended use and where performance may be unreliable or inappropriate. More nuanced language would help ensure that this provision preserves provider ability to maintain commercially necessary safeguards, system boundaries, and intended-use limitations while still supporting lawful agency use within the scope of the procured service.

Additional clarity is also needed regarding the scope of products covered by the clause itself. GSA has aligned the definition of "Artificial Intelligence System" with Section 7223(4)(B) of the *Advancing American AI Act*, which excludes common commercial products within which artificial intelligence is embedded, such as word processors, navigation systems, and other widely used digital tools where AI is not the primary purpose of the product. To avoid unnecessary uncertainty in implementation, GSA should explicitly state that common commercial products are not subject to these terms and conditions unless their primary commercial use is AI and the principal purpose of the agency's procurement is to acquire and use the product's AI capabilities. Without this clarification, the clause could unintentionally sweep in a broad range of commercial software products that incorporate AI-enabled features but are not procured as AI systems in any meaningful acquisition sense. Clarifying applicability in this way would better align the clause with existing statutory definitions while ensuring that compliance obligations remain focused on products where AI functionality is central to contract performance.

American AI Systems

TechNet supports policies that strengthen U.S. leadership in AI and ensure that federal procurement prioritizes trusted technologies developed by companies aligned with U.S. national and economic interests. Current federal policy appropriately encourages agencies to maximize the use of American-made AI and reflects a broader objective of strengthening domestic technological leadership while preserving access to commercially available innovation. The proposed GSAR clause, however, goes

materially beyond that policy direction by requiring that only “American AI Systems” be used in contract performance and by extending that requirement to any AI components manufactured, developed, or controlled by non-U.S. entities. As drafted, this creates a far more restrictive standard than current federal guidance outlined in OMB M-25-22, which outlines the American AI preference as a policy direction to ‘maximize’ use of American AI, not an absolute prohibition on any component with any foreign nexus. The clause’s categorical prohibition goes significantly further than M-25-22 requires and is likely inconsistent with its broader goal of ensuring the government has access to the best available AI capabilities. In so doing, it introduces substantial uncertainty into how agencies and contractors would determine compliance in modern commercial AI environments. It risks sweeping in a substantial share of the commercial AI offerings currently available to federal agencies, including products developed and led by U.S.-based companies and foreign companies (not from countries of concern) with U.S. subsidiaries that employ tens of thousands of Americans that rely on open-source components, globally distributed infrastructure, multinational research collaboration, or layered service delivery across commercial partners.

That broad restriction is also difficult to reconcile with President Trump’s Executive Order 14271 on “Ensuring Commercial, Cost-Effective Solutions in Federal Contracts,” which directs agencies to prioritize commercially available solutions and avoid imposing government-unique requirements that unnecessarily reduce competition or increase procurement costs. Modern AI systems are built through highly interconnected technology supply chains that do not fit neatly within a binary domestic-versus-foreign framework. Many leading U.S. AI providers incorporate open-source code, rely on globally maintained software libraries, use infrastructure components sourced through international supply chains, and partner across the value chain with cloud providers, infrastructure vendors, model developers, and application-layer providers operating in multiple jurisdictions. A requirement drafted this broadly risks excluding commercially available offerings developed by U.S.-based companies simply because they rely on widely used technical components that reflect how modern software ecosystems function.

The “American AI Systems” requirement has particularly significant implications for nearly every company that offers AI through modern cloud environments, model marketplaces, or interoperable service platforms. Many providers today offer customers access to multiple models, open-source options, retrieval systems, and integrated tools precisely to give agencies flexibility, avoid vendor lock-in, and support mission-specific deployment choices. A broad interpretation of “AI components” could unintentionally exclude offerings that include commercially standard open-source elements or model interoperability features that are now foundational to how agencies access AI securely and efficiently.

The provision also creates uncertainty for companies that work across allied technology ecosystems. U.S. AI leadership increasingly depends on trusted collaboration across democratic technology markets, including infrastructure partnerships, research collaboration, and software integration with companies

operating in allied countries. A domestic requirement that is not carefully tailored risks disrupting those relationships without advancing clear national security objectives.

If the policy objective is to prevent procurement of AI systems developed by foreign adversaries or models that present identifiable national security concerns, the clause should address that objective directly. A more targeted approach would focus on restricting procurement of AI systems developed, controlled, hosted, or materially influenced by prohibited foreign adversary entities, rather than adopting an overly broad requirement that creates sweeping uncertainty across the commercial market. Clarifying that distinction would better support both national security and continued American AI leadership while preserving competition and agency access to secure commercial technologies.

Conclusion

The federal government has an important opportunity to establish an acquisition framework for artificial intelligence that strengthens security, supports agency adoption, and expands access to the most advanced commercial technologies available. Achieving that objective requires procurement terms that reflect how modern AI systems are developed, licensed, secured, and continuously updated across the commercial market. When acquisition requirements depart too far from established commercial practice, the result is often reduced participation by leading providers, diminished competition, and slower government access to innovation.

To support that objective, TechNet respectfully urges GSA to revise the draft clause before incorporation into MAS Refresh 31 to ensure that any final requirements align with established commercial acquisition policy, existing government-wide security frameworks, and the operational realities of commercial AI deployment.

At a minimum, GSA should:

- Narrow provisions governing intellectual property ownership, licensing, and custom development to preserve commercially standard rights in proprietary systems and derivative technologies.
- Align incident reporting, documentation, and security obligations with existing FedRAMP and Security Requirements Guide processes rather than creating unclear or parallel compliance structures.
- Revise subjective performance provisions that introduce undefined standards and duplicative liability.
- Tailor change management requirements so they reflect the continuous update cycles that define modern AI services.
- Revise the “American AI Systems” requirement so that it advances trusted procurement objectives without creating unintended barriers to competition or excluding commercially available technologies that rely on modern supply chains, open-source components, and interoperable service models.

The success of federal AI procurement will depend on whether acquisition policy can preserve the speed, flexibility, and commercial scalability that define today’s leading

AI services while still meeting legitimate government security and accountability objectives. TechNet appreciates the opportunity to provide these comments and stands ready to engage in further dialogue to develop workable solutions that protect government interests while preserving contractors' ability to deliver innovative, high-quality AI services at scale.

Sincerely,

A handwritten signature in blue ink that reads "Linda Moore". The signature is written in a cursive, flowing style.

Linda Moore
President and CEO