

COUNSEL COMMENTARY

GSA Rolls Out New Standard AI Clause

A draft contract clause would impose sweeping new requirements on contractors and their AI model providers.

BY STEPHEN L. BACON

Earlier this year, the Department of Defense (DoD) and Anthropic, one of the nation's leading artificial intelligence (AI) companies, reached an impasse over the terms under which the DoD could use Anthropic's AI model, Claude. The DoD demanded the contractual right to use Claude for any lawful purpose, while Anthropic insisted on restrictions for applications related to lethal autonomous warfare and domestic mass surveillance.

Ultimately, the DoD designated Anthropic as a "supply chain risk" and banned it from federal contracting because it would not agree to provide Claude to the DoD on its preferred terms. Anthropic quickly challenged the legal validity of the DoD designation, which is the subject of ongoing litigation.

On March 6, 2026, within days of the DoD's "supply chain risk" designation against Anthropic, the General Services Administration (GSA) proposed a new clause, GSAR 552.239-7001, *Basic Safeguarding of Artificial Intelligence Systems* (proposed clause).¹ In addition to establishing that the government has the broad right to use AI systems "for any lawful government purpose," the proposed clause includes a range of new requirements that govern GSA contracts for AI capabilities. The proposed clause builds on and implements directives and policies set forth in Office of Management and Budget (OMB) Memorandum M-25-22, *Driving Efficient Acquisition of Artificial Intelligence in Government*.²

GSA requested comments on the proposed clause by April 3, 2026, offering a short window for contractors to provide feedback. Although the proposed clause would only apply to GSA solicitations and contracts, it will likely serve as a template for other federal agencies that may soon establish their own standard AI clauses. It is therefore important for all contractors and AI model providers to understand the potential ramifications of the proposed clause.

Coverage of the Proposed Clause

The proposed clause applies to all "solicitations and contracts for artificial intelligence capabilities." But the phrase "artificial intelligence capabilities" is not defined in the proposed clause and, thus, it is not clear how broadly it will apply.

However, certain uses of AI are excluded from the scope of acquisitions covered by OMB Memorandum M-25-22. Specifically, M-25-22 exempts the acquisition of "common commercial products within which artificial intelligence is embedded, such as a word processor or map navigation system." M-25-22 also excludes "AI used incidentally by a contractor during performance of a contract," meaning AI used at the contractor's option when not directed or required to fulfill contract requirements.

GSA should confirm that the M-25-22 carve-outs apply with equal force to the proposed clause, or clarify when a solicitation or contract is for "AI capabilities,"

such that it will be subject to the proposed clause. Absent additional clarity, contracting officers are likely to make inconsistent determinations regarding application of the proposed clause.

The "American AI Systems" Requirement

OMB Memorandum M-25-22 directed agencies to "maximize the use of AI products and services that are developed and produced in the United States." Although M-25-22 establishes a strong policy preference for American AI systems, the proposed clause imposes an outright ban on the use of any foreign AI systems and "any AI components manufactured, developed, or controlled by non-U.S. entities."

The proposed clause does not define the term "AI component," and it appears to impose a domestic sourcing requirement more stringent than existing standards governing federal software procurement. The Trade Agreements Act (TAA), for example, applies a "substantial transformation" test to determine a product's country of origin.

Under this test, a software product is generally "substantially transformed" in the location where the final software build takes place, regardless of where individual components or source code contributions originated.³ The proposed clause's new component-level origin test would appear to require a tracing exercise that the TAA has never demanded of software vendors.

In addition, the proposed clause inserts a “control” test for AI components that departs from the TAA’s focus on where “substantial transformation” occurs. AI components could potentially be deemed under the control of non-U.S. entities that have foreign investors, foreign board representation, or foreign development team members. GSA should either eliminate the “controlled by” standard, or define how it will determine whether an entity is subject to foreign control.

Contractor Responsibility for AI Service Providers

The proposed clause defines a “service provider” to mean “an entity that directly or indirectly provides, operates, or licenses an AI system but is not a party to the contract.” The proposed clause is clear that a contractor that uses an AI system of a service provider is responsible for that service provider’s adherence to the terms of the proposed clause.

To the extent there is any conflict between the proposed clause and an agreement between a contractor and a service provider, the proposed clause expressly states that it will govern. Thus, contractors must require their service providers to adhere to the requirements of the proposed clause even if they differ from the service provider’s standard commercial terms.

For example, as noted above, the proposed clause specifies that the contractor grants the government a license “to use the AI system for the duration of this contract for any lawful government purpose.” Thus, any usage restrictions in a service provider’s standard terms would not apply if they conflict with the government’s right to use the AI system for any lawful purpose.

Government Data Ownership and Use Restrictions

The proposed clause vests the government with full ownership of all “government data,” defined to include both data inputs and data outputs. Data inputs encompass everything submitted to the AI system, including all user prompts, queries, instructions, and documents. Data outputs are equally broad, covering everything generated by the AI system in performance of the contract including responses, analyses, data, and logs.

The prohibited use provisions in the proposed clause are also expansive. Contractors and their AI service providers may not use government data to train, fine-tune, or improve any AI model, for any customer or any commercial or non-commercial purposes. They also may not retain, access, or use government data beyond the scope and duration permitted by the contract.

The proposed clause mandates “eyes off” data handling procedures that restrict human review of government data except as strictly necessary to provide the AI system or respond to security incidents. All human access must be logged, justified, and limited to the minimum necessary for system functionality. Government data must be logically segregated from the data of non-government customers.

For companies whose commercial operations depend on using interaction data to improve their models, including for safety purposes, these restrictions may demand significant changes to the architecture and training of their AI systems.

Intellectual Property in Custom Developments

The proposed clause grants the government full, exclusive ownership of all custom developments, which are defined broadly to

include any modifications, configurations, enhancements, or fine-tuning performed specifically for the government under the contract. Custom developments must be dedicated to the government’s exclusive use, and contractors cannot use, reproduce, or derive benefit from custom developments without express written authorization from the contracting officer.

The proposed clause does specify that custom developments do not include background intellectual property (IP) developed independently by the contractor or service provider without use of or reference to the government’s information during the contract. In practice, however, it may be difficult to clearly distinguish existing models developed with background IP from improvements made under a government contract.

Fine-tuning a model on government data does not produce a distinct new artifact. Instead, the fine-tuning modifies the base model’s parameters in ways that may not be cleanly attributed to a particular source. This practical reality means that the background IP carve-out may be difficult to administer, and contractors should expect the government to broadly assert ownership over model improvements that the contractor may reasonably view as its own.

Unbiased AI Requirements and Government Evaluation Rights

The proposed clause establishes certain “Unbiased AI Principles” that contractors must make “commercial efforts” to satisfy. Specifically, the AI system must “be truthful in responding to user prompts seeking factual information or analysis,” and act as “a neutral, nonpartisan tool that does not manipulate responses in favor of ideological dogmas such as diversity, equity, inclusion.”

The government may conduct automated assessments of the AI system at any time using its own benchmarks, and is under no obligation to disclose those benchmarks, the underlying test data, or its evaluation methodologies to the contractor. The consequences of noncompliance are severe. The government may “suspend use of the AI System, until performance issues are satisfactorily addressed,” and if the contract is terminated for cause based on a violation of Unbiased AI Principles, the contractor is liable for the government’s “reasonable decommissioning costs.”

The proposed clause does not provide any definition of “performance issues” or “decommissioning costs.” The combination of vague compliance standards, undisclosed benchmarks, and potentially expansive liability creates a significant risk for contractors.

The uncertainty for contractors and service providers is compounded by another provision that provides that the “AI system must not refuse to produce data outputs or conduct analyses based on the contractor’s or service provider’s discretionary policies.” That requirement is difficult to reconcile with the proposed clause’s acknowledgment that compliance does not require retraining the model or altering model weights. For most commercial AI systems, content policies and safety guardrails are implemented at the model weight level and cannot be selectively disabled without retraining.

Disclosure, Reporting, Data Portability, and Change Management

In addition to the requirements outlined above, the proposed clause imposes

obligations that will require contractors to build new compliance processes:

- **Disclosure.** Contractors must identify all AI systems used in contract performance to the ordering contracting officer within 30 days of award, including whether any AI system has been modified to comply with a non-U.S. or commercial regulatory framework.
- **Incident Reporting.** Contractors must notify CISA and the contracting officer within 72 hours of any confirmed or suspected security incident, provide daily status updates until the incident is resolved, and preserve relevant logs and forensic artifacts for a minimum of 90 calendar days. Contractors must follow FedRAMP incident communication and response procedures to the extent they conflict with the proposed clause.
- **Data Portability.** Contractors must use open, standard data formats and application programming interfaces (APIs) for all data outputs and must provide export tools that allow the government to extract all government data, including conversational history, uploaded documents, and custom knowledge bases, in machine-readable formats sufficient for reconstruction in a separate system. These provisions implement M-25-22’s goal of preventing vendor lock-in situations.
- **Change Management.** Contractors must provide 30 days’ advance access to any successor model before discontinuing an existing one. In addition, contractors must notify the government within seven

days of identifying any change that materially increases bias or decreases safety guardrails.

Conclusion

The federal government is moving quickly to establish formal standards for AI procurement, and GSA’s proposed clause represents a significant step in that direction. Once finalized, GSAR 552.239-7001 will be considered for implementation in Multiple Award Schedule Solicitation Refresh 32 and other new contracts for AI capabilities issued by GSA. Contractors and AI model providers should expect that similar standard terms applicable to non-GSA contracts will be established in the near future. **CM**

Stephen L. Bacon is a shareholder in the Washington, D.C. office of the law firm Rogers Joseph O’Donnell, where he represents government contractors in bid protests, claims, terminations, investigations, and suspension and debarment proceedings. He frequently litigates cases at the U.S. Court of Federal Claims, the Government Accountability Office, the Boards of Contract Appeals, and the Small Business Administration’s Office of Hearings and Appeals. He also provides advice and counseling to clients on a broad range of contractual and regulatory compliance issues that confront government contractors.

The views expressed in this article are those of the author and do not necessarily reflect the views of Rogers Joseph O’Donnell or its clients. This article is for general information purposes and is not intended to be and should not be construed as legal advice.

ENDNOTES

- 1 GSA MAS Solicitation Refresh 31 Advanced Notice (Mar. 6, 2026) (“Advanced Notice”). The draft GSAR 552.239-7001 and its prescription regulation, GSAR 539.71, are included as attachments to the Advanced Notice.
- 2 OMB M-25-22, Apr. 3, 2025.
- 3 U.S. Customs and Border Protection, HQ H192146, Jun. 8, 2012; U.S. Customs and Border Protection, HQ H268858, Feb. 12, 2016; U.S. Customs and Border Protection, HQ H301776, Aug. 7, 2019.