August 3, 2026

ACT Association for Competitive Technology

General Services Administration
1800 F Street NW
Washington, District of Columbia 20405

RE: Public Comments of the Association for Competitive Technology (ACT) on Proposed GSAR Clause 552.239-7001, Basic Safeguarding of Data Within Large Language Model Artificial Intelligence Systems | Notice-MVAC-2026-01; Docket No. GSA-2026-0331

In response to the request for comments published on June 17, 2026,[1] the Association for Competitive Technology (ACT) submits these comments to the General Services Administration (GSA). ACT appreciates the opportunity to provide feedback on the draft of new General Services Administration Acquisition Regulation (GSAR) clauses (48 C.F.R. Parts 539 and 552) regarding the government's procurement of large language models (LLM or LLMs). ACT respectfully submits these comments to provide input to the GSA as it tackles the difficult problem of procuring cutting-edge and performant LLMs across the federal government.

ACT welcomes the opportunity to comment and shares the GSA's goals of protecting government data while also promoting American artificial intelligence (AI) innovation. We offer these comments from the perspective of the small business technology community, which depends on clear and predictable contractual terms and practices as they develop for and compete in markets being shaped and powered by LLMs. ACT believes that any final GSAR should disclose how contractors' LLMs will be evaluated against the proposed Unbiased AI Principles, clarify that contractors do not need to extend flowdown requirements to open-source developers who do not host services, and allow contractors to evaluate risks using the National Institute of Standards and Technology's (NIST) Risk Management Framework.[2]

ACT's principal recommendations, developed in the comments that follow, are:

- Disclose, at contract time, a high-level and immutable description of the objective benchmarks the government will use to evaluate LLMs against the Unbiased AI

[1] General Services Administration, Request for Comments on the Acquisition of Information and Communication Technology, 91 Fed. Reg. 36559 (June 17, 2026), Docket No. GSA-2026-0331.
[2] NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY, NIST SPECIAL PUBLICATION 800-37 REVISION 2 R, RISK MANAGEMENT FRAMEWORK FOR INFORMATION SYSTEMS AND ORGANIZATIONS (2018) https://doi.org/10.6028/NIST.SP.800-37r2 [NIST RISK MANAGEMENT FRAMEWORK].

Principles, and provide a defined, reasonable cure period before any termination for non-compliance.
- Clarify the Unbiased AI Principles so they do not sweep in ordinary, good-faith training-data decisions inherent to building any LLM.
- Narrow and clarify the clause's applicability by tying it to the amount and sensitivity of the government data an LLM will process.
- Replace absolute "material change," foreign-component, and "no security risk" language with a risk-based standard modeled on the NIST Risk Management Framework.
- Exclude open-source and open-weight developers who do not host services from the LLM-role definitions and flowdown obligations.
- Preserve standard commercial terms by limiting flowdown to reasonable and practicable steps, rather than overriding contractors' commercial agreements.
- Recognize the disproportionate burden on small, micro, and non-technology businesses; harmonize the clause with existing federal AI guidance; and pair any final clause with small-business technical assistance and a small-entity impact analysis.

I. Introduction and Statement of Interest

ACT is a trade association that represents small and medium-sized technology companies that develop software powering the modern digital economy. Our members are entrepreneurs, innovators, and independent developers within the global app ecosystem who build products and services across every industry, including the AI and LLM industry. ACT represents a domestic ecosystem valued at approximately $1.8 trillion, supporting 6.1 million American jobs.[3] We work with our members and policymakers to promote a policy environment that rewards and inspires innovation while providing resources to help developers create jobs and continue to build technology.

As smaller players within a large ecosystem, the freedom and ability of larger companies to incorporate ACT members' software into their own solutions is critical as is the clarity and predictability of government contracts. The small business technology developers ACT represents provide software to government contractors both as service providers and as open-source software developers. ACT's members additionally develop specialized, domain-specific LLMs. We therefore approach this proposed draft of a new GSAR procurement clause with two principles in mind:

- Any GSAR clauses must be clear and practicable regardless of whether they affect prime contractors or subcontractors; and
- Any government contracts must be predictable and stable so that ACT's members can rely on these contracts to invest significant capital and personnel resources to

[3] Association for Competitive Technology (ACT), *State of the App Economy*, available at https://actonline.org/wp-content/uploads/APP-Economy-Report-FINAL-1.pdf.

develop software to support government contractors and compete for government contracts themselves.

II. General feedback on the draft of the proposed GSAR clause

ACT is thankful for the thoughtful improvements and community engagement the GSA made between the proposed draft of this GSAR clause announced on March 6, 2026, through a GSA Schedule refresh and the current draft of this GSAR clause. ACT appreciates that the scope of the draft clause is narrowed to exclude common commercial products or products where the LLM functionality is incidental to the purpose of the government contract. ACT additionally is glad that the GSA has clarified that companies retain the intellectual property of the training data and model weights companies develop prior to and independent from any government contracts.

In the spirit of constructive collaboration, ACT offers general comments on the draft of language of the proposed GSAR clause.

1. The evaluation of contractors' adherence to "Unbiased AI Principles" leaves companies vulnerable to their contracts being arbitrarily terminated.

Like many commenters at the GSA's July 14th listening session, ACT is deeply concerned that the methods used to evaluate LLMs' adherence to the GSAR's "Unbiased AI Principles" are not required to be disclosed to contractors. ACT believes the undefined evaluation procedures and contractors' responsibility for decommissioning costs will cause government contracts to be perceived as unstable. This perception will depress small businesses' ability to compete with established companies and also decrease small businesses' willingness to supply software services to contractors seeking government LLM contracts.

i. *The government not being required to disclose its evaluation methods increases the cost of government contracts.*

On page 36565 proposed sections (j)(2)(iv) and (j)(3)(ii) state that the government "is under no obligation to disclose or provide access to the underlying data [or] methodologies" used to evaluate the performance of an LLM except when those data or methodologies are used as the basis to "terminate" a contract for "failure to remediate . . . non-compliance . . . with the Unbiased AI Principles." Under section (j)(3)(ii)(A), contractors are responsible for decommissioning costs if the government terminates a contract for a failure to adhere to the Unbiased AI Principles.

ACT shares the GSA's goals of keeping government data secure and critical infrastructure stable but is deeply concerned that companies will not have access to the methodologies used to evaluate their LLMs until the government is already pursuing termination of their contract. ACT also acknowledges the government's desire to disclose specific tests so that

companies don't build their LLMs to only pass a specific test. Still, because the government is not required to share any of the data or methodology—including even a high-level description of those tests—used to evaluate the performance of an LLM, companies will need to build and maintain LLMs without any way to determine if their LLMs are performing as expected.

Contractors could spend significant time and effort ensuring their LLMs adhered to the government's Unbiased AI Principles and still have their contracts terminated if their reasonable interpretation of the principles differed from the government's. Further still, agencies' interpretation of these principles could vary, creating uncertainty for each contract a contractor implements. The uncertainty sections (j)(2)(iv) and (j)(3)(ii) introduce will negatively impact government procurement as companies raise the price of their contracts to account for the significant capital and personnel costs they would expend if they need to remediate and possibly decommission any previously undefined deviation from the government's Unbiased AI Principles.

Furthermore, under section (j)(3)(ii), there is no minimum defined timeline between when the government notifies the prime contractor of non-conformance and when the government can terminate the contract for a "failure to remediate" an LLM. Under this language, companies could receive notice about an LLM's non-compliance with a standard and not be provided an appropriate amount of time needed to correct that non-compliance. Again, this uncertainty harms the government as prime contractors will need to evaluate the likelihood of their contracts suddenly being terminated and price in any decommission costs they are responsible for under section (j)(3)(ii)(A).

Taken together, sections (j)(2)(iv) and (j)(3)(ii) subject contractors to the possibility that the methods used to evaluate the performance of their LLMs can arbitrarily change causing their contracts to be terminated capriciously. Given the background of the Department of War's dispute with Anthropic, these arbitrary decisions are not just theoretical.[4] The methodologies used to evaluate the performance of a contractor's LLM could change when government administrations and agency heads change or if a government official was unhappy with how a company negotiated a government contract.

Again, the government's undisclosed methods used to evaluate LLMs increases the cost of government contracts as companies incorporate the risk of these contracts. Additionally, sections (j)(2)(iv) and (j)(3)(ii) together decrease competition in the LLM industry as small businesses often lack the robust legal departments needed to interpret vague contract requirements. Further still, the quality of LLMs the government is able to purchase could decrease if companies choose not to contract with the government because they are

[4] Anthropic, PBC v. Department of War, No. 26-1049 (Fed. Cir. 2026). ACT has appeared as amicus curiae in that matter. See Brief of the Association for Competitive Technology as Amicus Curiae Supporting Petitioner at 13–14, Anthropic, PBC, v. Department of War et. al, No. 26-1049 (Fed. Cir. Mar. 17, 2026) https://actonline.org/wp-content/uploads/17-Mar-2026-ACT-Amici-Motion-for-Leave-Amicus-STAMPED.pdf.

unwilling to enter contracts that could be terminated suddenly and without time to correct any non-compliance.

ii. The definition of Unbiased AI Principles is unworkable as written and should be clarified.

In addition to the uncertainties regarding the evaluation of Unbiased AI Principles, that definition, as written, is unworkable because it would place virtually all LLMs in violation of those principles. On page 36565, section (j)(1)(ii), contractors must not "intentionally introduce or embed partisan or ideological judgements into the LLM[] . . . through methods such as training data selection . . . ."

Training any LLM involves some degree of intentional training data selection.[5] As there is an increasingly infinite amount of written text available in print and on the internet, LLM developers must decide what training data to use to fit the computational constraints inherent in any computer system. Thus, section (j)(1)(ii)'s requirement not to "intentionally" introduce bias into an LLM is fundamentally at odds with how LLMs are built because all LLM developers are required to choose which out of an infinite amount of data sources they should include in their training data.

Furthermore, section (j)(1)(ii)'s requirements about not intentionally introducing ideological bias into an LLM prevents contractors from using open source, open-weight models to cost-effectively build domain-specific LLMs. LLMs are developed by distilling the patterns from a vast amount of data into numerical "weights."[6] Developers can publicly share the numerical "weights" of their LLM as an open-weight model.[7] Instead of building a domain-specific LLM from scratch, others can efficiently and cost-effectively build on top of these open-weight models because they only need to tweak the models to perform in a specialized domain.[8]

---

[5] Cole Stryker, *What are Large Language Models (LLMs)?,* IBM https://www.ibm.com/think/topics/large-language-models ("Large language models (LLMs) are a category of deep learning models trained on immense amounts of data . . . ."); AI Demystified: Introduction to Large Language Models, STANFORD UNIVERSITY IT, (Dec. 13, 2024) https://uit.stanford.edu/service/techtraining/ai-demystified/llm ("Training Data: LLMs are trained on extensive [but non-exhaustive] datasets . . . ."); *What is an LLM (Large Language Model)?*, AMAZON WEB SERVICES, https://aws.amazon.com/what-is/large-language-model/ ("Large language models . . . are pre-trained on vast amounts of data.").

[6] Kate Conger, Mike Isaac, & Meghan Tobin, *What is Open-Weights A.I.?*, N.Y. TIMES, (July 28, 2026) https://www.nytimes.com/2026/07/28/technology/open-weight-ai.html.

[7] *Id.*

[8] Haiman Wong, *Mapping the Open-Source AI Debate: Cybersecurity Implications and Policy Priorities*, R STREET, (Apr. 17, 2025) https://www.rstreet.org/research/mapping-the-open-source-ai-debate-cybersecurity-implications-and-policy-priorities/#_edn45 ("By participating in open-source initiatives, technology firms . . . benefit from [] collective innovation . . while reducing the costs of [AI] research and development." (highlighting Cailean Osborne, *Why Companies 'Democratise' Artificial Intelligence: The Case of Open Source Software Donations*, arXiv, (Sept. 26, 2024) https://arxiv.org/html/2409.17876v1).

Contractors lack the capability to audit the training decisions a third party used to develop an open-weight LLM. Despite their usefulness, contractors are effectively barred from using open-weight models because they cannot certify nor prove that *no* ideological biases were intentionally introduced into an LLM built on top of an open-weight model. Section (j)(1)(ii) stifles American AI innovation because the small businesses with expertise in specialized domains, ACT's members, lack the capital to build LLMs from scratch which section (j)(1)(ii) effectively requires. This, furthermore, harms the government because they cannot realize the cost-efficiencies of developing specialized models on top of open weight models.

* * *

To encourage American AI innovation and enable the government to purchase performant and cost-effective LLMs, the GSA should amend how the performance of LLMs is evaluated. First, the GSA should amend section (j)(2)(iv) to require the government to share the methods used to assess the performance of an LLM. At minimum, the government should be required to share, at contract time, a high-level and immutable description of the features they will use to evaluate LLM performance (which should be objective and benchmarked as much as practicable so that contractors of all sizes can assess compliance in advance). Second, the GSA should amend section (j)(3)(ii) to give prime contractors a reasonable amount of time to remediate LLM non-compliance. And third, the GSA should change section (j)(1)'s definition of "intentionally" introducing bias into LLMs to better reflect the reality of LLM development.

2. ***The broad applicability of the LLM Procurement clause 552.239-7001 has improved but needs additional clarification.***

ACT appreciates that the GSA narrowed the applicability of the LLM Procurement clause 552.239-7001 but is concerned that the updated clause is still unclear. On page 36561, section (a)(1) states that the clause applies "only when Government Data will be processed" by an "LLM" except when the LLM is "embedded in a common commercial product" or when the "LLM functionality is incidental to the primary purpose of the core requirement being procured." Deciding if a potential government contract is subject to this new procurement clause is critical because, under 539.71 on page 36561, Contracting Officers "must" include this clause when section (a)(1) applies.

It is unclear how section (a)(1) applies when companies add LLM features to existing government contracts. Companies iteratively develop and release software by adding new features to existing applications. Depending on a new feature's popularity, companies may expand that feature and incorporate it into core product workflows. It is unclear if clause (a)(1) would apply to a situation in which a company entered a government contract without an LLM embedded in their application and then later added an LLM that became part of the application's core functionality. Initially, clause (a)(1) would not apply because there was no LLM in the product when the contract was executed. Later, however, if the

LLM was frequently used but no longer incidental to the government's "purpose" of the contract, it is unclear if clause (a)(1) would apply during the contract and if the contract would be extended.

Contrary to the GSA's expectations, section (a)(1)(ii)(A) might not apply to a contract whose primary purpose is procuring an LLM if that LLM is popular enough to be classified as a "common commercial product." Many companies have already deployed LLM systems to their employees.[9] As companies continue to deploy LLMs, specific LLMs could be considered "common commercial product[s]." Then, evaluating whether clause (a)(1)(ii)(A) applies would turn on whether that LLM is considered a "common commercial product." That could generate an odd situation where well-known, common LLMs, such as OpenAI's ChatGPT, would not be subject to clause 552.239-7001 but domain-specific LLMs, such as health- or finance-focused LLMs, would be subject to 552.239-7001. This would stifle innovation because larger, more common LLMs would not need to adhere to the additional requirements of the proposed GSAR clause whereas developing, non-commercialized, domain-specific LLMs would be.

To make the application of the clause (a)(1) clearer and easier to evaluate, ACT recommends the GSA amend the definition to depend on the amounts and types of government data intended to be entered into an LLM.

3. The definition of "material change" is overbroad.

As written, the definition of "material change" is unclear and overbroad because any software changes *could* affect the security or performance of a contract. On page 36562 under section (b), the definition of "material change" includes "any modification that could affect the trustworthiness, security, or operational integrity" of the contract. This definition of "material change" is also flowed down to all subcontractors on page 36566.

Even if that degree of risk is low, any software change carries the risk that it *could* negatively impact a system no matter how small, routine, or thoroughly tested.[10] For example, consider the following scenarios where common-place and low-risk changes caused system outages and failures.

- Routine updates of system packages introduced security vulnerabilities.[11]

---

[9] Jeffery S. Allen, *Monitoring AI Adoption in the US Economy*, FEDERAL RESERVE (Apr. 3, 2026) https://www.federalreserve.gov/econres/notes/feds-notes/monitoring-ai-adoption-in-the-u-s-economy-20260403.html ("[A] November iteration of the Survey of Business Uncertainty . . . estimates that . . . about 54 percent [of the labor force] works at [U.S.] firms that use LLMs.").

[10] NIST RISK MANAGEMENT FRAMEWORK, *supra* note 2, at 43 ("Protection needs . . . [convey the system requirements needed] to reduce security and privacy risk to an acceptable level . . . .").

[11] *See* Microsoft Defender Security Researcher Team, *Malicious npm Packages Abuse Dependency Confusion to Profile Developer Environments*, (May 29, 2026), https://www.microsoft.com/en-us/security/blog/2026/05/29/33-malicious-npm-packages-abuse-dependency-confusion-profile-developer-environments/.

- Commands approved to fix non-urgent system issues caused an outage when they were manually mistyped.[12]
- Changing the text of a user-facing message caused a failure.[13]

Modern software systems are increasingly made of multiple systems which can each be updated many times a day. First, today's software systems often contain multiple components each with a separate responsibility, such as processing user requests, detecting fraudulent transactions, authenticating users, storing user actions to an audit log, and maintaining the integrity of data storage. Each one of these components can be updated by human decisions multiple times a day through code changes, configuration changes, or by system administrators taking actions to remediate system irregularities. Additionally, automated tools can change these components multiple times per *hour;* common infrastructure management tools can alter the resources allocated to software systems to address fluctuations in the volume of user requests.[14] Together, these human-initiated and automated changes are numerous and any one of them "could" affect the performance of an LLM system.

Consolidating and reporting these changes is technically infeasible. Prime contractors would need to collect and consolidate the hundreds, if not thousands, of human-initiated and automated daily changes made to their system. As subcontractors' changes "could" also impact the performance of an LLM system, subcontractors would also need to collect all of their changes that "could" impact the prime contractor. Collecting and reporting the numerous changes would place a technically infeasible burden on prime contractors and subcontractors alike.

Additionally, reviewing the large volume of reports of *any* change that *could* affect the integrity of a system would not be useful. To detect any changes which could likely affect the performance of a contract, the Contracting Officer would need to review many changes a day without the knowledge of the architecture of the system they were analyzing. Furthermore, reviewing the summary of these changes would provide the government with an incomplete overview of changes that could negatively impact the performance of their LLM because many vulnerabilities are latent and could be introduced weeks, months, or even years before they are discovered.[15] A government Contracting Officer reviewing the

---

[12] *See Summary of the Amazon S3 Service Disruption in the Northern Virginia (US-EAST-1) Region*, *https://aws.amazon.com/message/41926/*

[13] *Cisco Bug: CSCsr01315*, CISCO, (May 25, 2026), https://bst.cisco.com/quickview/bug/CSCsr01315 ("Change in the Error Message causes failure in Release-4.")

[14] *See* Adrien Payon & Shaoni Mukherjee, *A Deep Drive into Cloud Auto Scaling Technologies*, DIGITAL OCEAN, (July 31, 2026) https://www.digitalocean.com/community/tutorials/auto-scaling-techniques-guide

[15] *See, e.g., Fastly Says Single Customer Triggered Bug Behind Mass Internet Outage*, THE GUARDIAN, (June 8, 2021, 10:02 EDT), https://www.theguardian.com/technology/2021/jun/09/fastly-says-single-customer-triggered-bug-that-caused-mass-outage ("A bug in Fastly's code introduced in mid-May had lain dormant until [early-June] . . . ."); *OCI Status*, ORACLE CLOUD, (July 20, 2022) https://ocistatus.oraclecloud.com/#/incidents/ocid1.oraclecloudincident.oc1.phx.amaaaaaavwew44aa7zo

numerous changes that *could* impact the performance of a system, would be an inefficient use of time and resources and provide little value.

ACT recommends that the government amend the definition of "material change" to be "any change *likely* to affect the trustworthiness, security, or operational integrity of the contract as assessed by weighing the importance of the particular system changed against the degree of risk that change poses."[16] That updated definition would relieve the intractable reporting burdens from prime contractors and subcontractors; it would also save the government Contracting Officer time because contractors and subcontractors with the appropriate expertise needed to evaluate risk would assess and then only report relevant changes.

4. ***Section (i)(4) requiring contractors to report "a[n]" unreported "change" is overbroad.***

On page 36565, section (i)(1) only requires contractors to notify the Contracting Officer of "planned material changes." However, section (i)(4) requires contractors to notify the Contracting Officer of "a change" if the contractor implements said change "without providing sufficient advance notice." The definition of "a change" is not limited to be only changes made due to security or service emergencies because of where the "or" is placed in section (i)(4). These changes are flowed down to LLM System Developers, LLM System Operators, and LLM Service Providers on page 36566.

As written, section (i)(1) would require parties to notify the government Contracting Officer of planned *material* changes, yet section (i)(4) could be read to require parties to notify the Contracting Officer of *a* change made without notice. Although sections (i)(1) and (i)(4) are likely intended to be read together, due to "a change" being limited to changes without "sufficient advance notice," section (i)(4) does not explicitly reference section (i)(1). Thus, there is a plausible reading of section (i)(4) that requires parties to inform the Contracting Officer of *any* changes made at all.

When operating and deploying software, many *non-material* changes are made as part of the normal course of software development. These non-material changes can be made manually by software engineers deploying code changes or by automated systems. Collecting and reporting these routine and non-material changes is technically impracticable and could inundate the Contracting Officer with non-useful information.

skanlspjh4ll6wxhwxrbkbed4d4cnupxexzqzvlyq ("Following unseasonably high temperatures . . . two cooler units in the data center experienced a failure when they were required to operate above their design limits.").

[16] This aligns with NIST's Risk Management Framework. NIST RISK MANAGEMENT FRAMEWORK, *supra* note 2, at 38–45.

ACT recommends that the GSA amend section (i)(4) to only require parties to notify the government of *material* changes in addition to changes initiated to address security or service emergencies.

5. ***The absolute language about "not" introducing security risks through "foreign-developed" components is unrealistic and discourages companies from incorporating open-source software despite the GSAR clause explicitly allowing open-source software.***

ACT appreciates the GSA's goal of ensuring that U.S. government data cannot be compelled to be disclosed by adversarial foreign governments, but cautions that the language of section (f)(2)(iii) is too narrow to be operationalized.

On page 36566, proposed GSAR section (f)(2)(iii) permits "foreign-developed" components, e.g., "open-source components" and "globally operated infrastructure dependencies," so long as they "do not introduce security risk or foreign control" risk. Contractors must ensure that foreign-developed components "do not" introduce a risk that an LLM could be "controlled by a foreign government" and that any "core" LLM models are prohibited from being "developed" by "entities subject to the . . . influence . . . of adversary foreign governments."

Again, ACT shares the GSA's goals of ensuring that U.S. government data is protected from adversarial governments but wishes to highlight that requiring companies to only use and operate components in a way that does "not" pose any security risk is unrealistic. As discussed above, any software change or deployment includes some degree of risk, including security risks.[17] No matter how carefully developers write code, *all* software is at risk of being shipped with unknown vulnerabilities.[18] If adversarial governments discover those vulnerabilities before software developers can fix them, the same adversarial governments can exploit those vulnerabilities to gain access to and "control" an LLM. The GSAR's proposed clause of allowing entities to use foreign-developed components so long as they "do not" introduce a "security risk" is fundamentally flawed.

Likewise, the proposed GSAR language allowing contractors to use "open-source components" and "globally operated infrastructure" so long as they are "not" "developed" or "operated" by entities that could be "influence[d]" by an "adversarial foreign government" is fundamentally flawed. First, open-source projects are often developed by many contributors across the globe. It is impracticable to develop a list of people who "developed" any open-source software component in the vast software supply chain. It is

[17] NIST RISK MANAGEMENT FRAMEWORK, *supra* note 2, at 43 ("Protection needs . . . [convey the systems requirements needed] to reduce security and privacy risk to an acceptable level . . . .").
[18] Richard Ciancioso, Danvers Budhwa, Thaier Hayajneh, *A Framework for Zero Day Exploit Detection and Containment*, INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS (Apr. 2, 2018) ("All systems with software code are susceptible to an attack called "zero day", which exploits vulnerabilities that are not yet disclosed or known to the public or vendor.")

cost prohibitive[19] to audit numerous individuals' lives to see if they were subject to an iota of "influence" from an adversarial government.[20] Similarly, running "globally operated infrastructure" is subject to the same impracticalities. Contractors would need to get the names and information of any operator conceivably in the operational supply chain, then audit their lives to see if they could be "influence[d]" by an adversarial government. Despite the proposed GSAR's language allowing the use of "open-source components" and "globally operated infrastructure," the absolute security and foreign control language practically bars contractors from using them.

ACT recommends that the GSA amends the proposed clause to allow contractors to use a "risk-based approach," as stated in section (f)(2)(iii)(C), when evaluating foreign-developed and globally operated components. As described in NIST's Risk Management Framework, contractors should assess a component's likelihood and consequence of loss to determine if using foreign-developed or globally operated components is appropriate for a particular system.[21]

6. ***The definition of LLM Developer could be read to require an external security researcher to be subject to GSA's terms even if they don't work on an LLM that could access government data.***

ACT is concerned that the definition of LLM Developer discourages companies from using third-party security researchers or beta testers to assess the security of LLM models before they are released publicly. On page 36562, clause (b) defines "LLM Developer" to mean "the party" that performs "LLM Development Tasks [by] . . . testing . . . or otherwise making available an LLM . . . ." On pages 36565–66, LLM Developers are required to ensure the *Performance, Evaluation, and Remediation* requirements under clause (j).

The definition of LLM Developer can be read to include third-party security researchers employed by a contractor to "test" new LLMs or users who "test" a "beta" version of an LLM before it is released publicly. That interpretation would cover these third-party researchers and users even if the LLM they evaluated never processed government data. Furthermore, the responsibility to "ensure" the LLM is developed with the "unbiased AI principles" would flow down to these external testers through clause (j) despite their inability to exert control of the development of an LLM due to their limited scope of influence.

---

[19] MICHELLE D. CHRISTENSEN, CONG. RSCH. SERV., R43216, SECURITY CLEARANCE PROCESS: ANSWERS TO FREQUENTLY ASKED QUESTIONS 9 n.47 (2023) (noting that Tier 3 [Secret] investigations bill at $420 and Tier 5 [Top Secret] investigation is $5,410 (citing the Defense Counterintelligence and Security Agency's "Position Designation Tool")).

[20] Even Top Secret background checks may clear people who can be influenced to share data with foreign governments. *See, e.g.,* Alexander Mallin, Former NSA Employee Admits Trying to Sell Top Secret Info to Russia, ABC NEWS, (Oct. 23, 2023) https://abcnews.go.com/Politics/former-nsa-employee-admits-sell-top-secret-info/story?id=104234977.

[21] *See* NIST RISK MANAGEMENT FRAMEWORK, *supra* note 2, at 38–45.

As shown by Anthropic's decision to develop Claude Mythos only to decline to release it publicly due to cybersecurity concerns, assessing the security impact of LLMs before they are released is important for public safety and internet security.[22] ACT is concerned that the definition of LLM Developer and the responsibilities it entails would discourage third-party security researchers and beta testers from agreeing to test soon-to-be-released LLM models. That lack of testing would detrimentally affect the quality and security of the LLMs furnished through government contracts because the LLMs would not be as robust or secure as they could be. Additionally, this clause could harm public safety and internet security if a model was released with a security bug that could have been caught by external researchers.

To encourage the secure development of LLMs and to protect the independence of third-party security researchers, ACT recommends the GSA amend the definition of LLM Developers to exclude individuals who "test" LLMs that do not interact with government data.

7. Contractors' inability to audit government data leaves them unable to confirm that the government is adhering to the scope of their contract.

ACT appreciates that the GSA needs to keep government data confidential but is concerned that the proposed GSAR clause prohibits companies from inspecting government data on a limited basis to ensure the government adheres to the scope of their license.

Under section (e)(2) on page 36563, the government is granted a license that is "strictly" limited to the specific purpose and scope of work defined in the contract. However, under section (e)(1)(iii), the contractor is granted a "limited" license to use government data "solely" for an enumerated list of purposes. The section that best allows the contractor to audit the government's usage of their LLM is section (e)(1)(iii)(A) which allows contractors to "use" government data for "[p]erforming the specific requirements" of the contract.

Because there is no clause in section (e)(1)(iii) that allows contractors to audit government data, even in a limited and secure manner, companies lack the ability to ensure the government isn't exceeding the scope of their license. "Performing the specific requirements" of a contract in section (e)(1)(iii)(A) could arguably be stretched to include ensuring that the government operates within the limits of their license, as *auditing* a contract is tangentially related to *performing* a contract, but that is uncertain. That uncertainty leaves contractors without a clearly defined way to ensure their software isn't used in a way they did not agree to.

[22] Frontier Red Team, *Assessing Claude Mythos Preview's Cybersecurity Capabilities*, (Apr. 7, 2026), https://www.anthropic.com/research/mythos-preview.

To protect contractors' contractual rights, the GSA should amend section (e)(1)(iii)(A) to allow contractors to assess that the government is adhering to the scope of their license by allowing them to view government data for auditing purposes in a limited and secure manner.

8. ***There is ambiguity in the proposed contract between contractors' responsibility to preserve data after a security incident and their responsibility to delete data after a contract is terminated.***

ACT is concerned about an ambiguity in the draft GSAR clause that requires contractors and subcontractors to delete government data after a contract is terminated yet also requires contractors to preserve data for at least "90 days" after a security incident.

On page 36563, proposed clause (e)(4)(vii) requires contractors to "securely and permanently delete" all government data at the end of a government contract. On page 36564, clause (f)(5)(iii) requires the contractor to "preserve all relevant logs, forensic images, and incident artifacts for a minimum of 90 days from a security incident." Additionally, on page 36562, section (e)(1)(iii) grants contractors a "limited, revocable" license that lasts "for the duration" of the contract. All subcontractors are required to implement clauses (e) and either (f) or (f)(4) through the flowdowns on page 36566.

If a security incident occurs 90 days before a contract ends, sections (e)(4)(vii) and (e)(1)(iii) could be read to require contractors and subcontractors to *immediately* delete government data. That reading would then conflict with clause (f)(5)(iii)'s requirement to preserve data. This ambiguity places contractors in a difficult position because they run the risk of deleting government data necessary to investigate a security incident or they run the risk of preserving government data after their license to store that data is revoked.[23]

To avoid uncertainties about these sections' interactions, ACT suggests the GSA clarify which section takes precedence if a security incident occurs 90 days before the expiration of a contract.

[23] They may also face criminal penalties under 18 U.S.C. § 641.

III. Answers to GSA's specific Questions 3 and 4: "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?" and "Do you understand how to implement the flowdown clauses?"

1. It is unclear if open-source maintainers are included in the definitions of LLM Developer, LLM System Operator, LLM System Integrator, and LLM Service Provider.

ACT is concerned that the definitions of LLM Developer and LLM System Integrator could reasonably be interpreted to include the open-source developers critical to modern software development. These overly broad definitions would impede LLM software development because companies using these technologies lack the capability to require open-source maintainers to adhere to the requirements in the flowdown.

Most companies incorporate open-source software into their software components.[24] These open-source components are used in a variety of systems including critical software components and LLMs.[25]

When companies incorporate open-source software into their technologies, they accept "as-is" open-source licenses that typically disclaim any responsibility arising from companies' use of that software.[26] Thus, the contractors and subcontractors providing software services to the government do not have the capability to flowdown any GSAR requirements to these third-party developers.

The definitions of the LLM roles required to adhere to the flowdown clauses, section (b) on page 36562, can be read to include open-source technologies. First, an LLM Developer is the party that "train[s]" or "publish[es]" "an LLM." An open-source developer who published an open-weight model later incorporated into an LLM would meet this definition. Second, an LLM System Integrator is the party that "performs LLM Design and LLM Deployment tasks by . . . configuring . . . how an LLM System performs in a specific

---

[24] *Open Source Software Market Size, Share, Growth, and Industry Analysis, By Type (Program Development,WEB Development,Mobile Phone/Mobile Development,Development Tools,Others), By Application (Large Enterprise,SME), Regional Insights and Forecast to 2034*, MARKET REPORTS WORLD, (Dec. 12, 2025) https://www.marketreportsworld.com/market-reports/open-source-software-market-14722346 ("[A]mong sampled commercial codebases, 96 percent contained open source components. More than 70 percent of enterprises reported using open source in mission-critical systems. . . ."). The underlying figures are drawn from the Synopsys/Black Duck Open Source Security and Risk Analysis (OSSRA) Report. See Black Duck, 2024 Open Source Security and Risk Analysis (OSSRA) Report (2024), https://www.blackduck.com/resources/analyst-reports/open-source-security-risk-analysis.html (reporting that 96% of scanned codebases contained open source).

[25] *Id.;* Conger, *supra* note 6; *What is Linux?*, OPENSOURCE.COM, https://opensource.com/resources/linux.

[26] *See, e.g., The 3-Clause BSD License*, OPEN SOURCE INITIATIVE, https://opensource.org/license/bsd-3-clause; *The MIT License*, OPEN SOURCE INITIATIVE, https://opensource.org/license/mit.

deployment . . . .” Again, a developer publishing a specialized open-weight LLM model meets this definition. Additionally, a provider of open-source database software meets this definition because they “configur[ed]” interfaces LLM Service Providers use to store and retrieve LLM conversations.[27]

The ambiguities about whether these definitions apply to the open-source software providers that do not host services impedes software development. The companies who build and deploy LLMs as well as the companies who provide software services to LLM Developers will need to assess if their contracts allow them to incorporate common, open-source software into their systems. Companies would either need to spend an exorbitant amount of time reimplementing common, open-source software or spend significant legal resources arguing why the open-source components they used do not meet the GSAR’s LLM role definitions. These burdens would fall particularly hard on ACT’s small business members who, relative to larger companies, often lack the resources to evaluate these ambiguities.

Beyond the flowdown problem, treating open-source and open-weight contributors as regulated parties would chill innovation in the U.S. open-source ecosystem on which cost-effective, domain-specific AI depends. If contributing to an open-weight model or a widely used software component could expose a small U.S. developer to federal contract obligations it cannot satisfy, many will simply decline to publish or maintain such components. The result would be greater dependence on a small number of large, closed commercial models, which is precisely the concentration that undermines competition, resilience, and American leadership in open AI development. Indeed, the Administration’s own AI Action Plan identifies open-source and open-weight models as a strategic priority for American AI leadership, so a procurement clause that penalizes the small U.S. developers who contribute to those models works against the government’s own stated goals.[28]

ACT recommends the GSA clarify that the open-source developers that provide “as-is” software are excluded from the definitions of LLM Developer, LLM System Operator, LLM System Integrator, and LLM Service Provider. Clarifying that open-source maintainers who do not host services incorporated into LLM systems would decrease the ambiguity of the proposed GSAR clause which would decrease the cost of building LLM software for the government.

2. *The proposed LLM procurement contract takes precedence over standard commercial terms which puts ACT’s members in a vulnerable position.*

[27] *See, e.g.*, *PostgreSQL: About*, THE POSTGRESQL GLOBAL DEVELOPMENT GROUP, https://www.postgresql.org/about/.

[28] The White House, Winning the Race: America’s AI Action Plan (July 2025) (identifying open-source and open-weight AI models as a national strategic priority); see also Open Source Initiative, White House Releases AI Action Plan, Includes Open Source (July 24, 2025), https://opensource.org/blog/white-house-releases-ai-action-plan-includes-open-source.

ACT is concerned that the proposed GSAR section (c), on page 36562, establishes requirements that "take precedence over [contractors'] conflicting" "terms, conditions, or commercial agreements." As smaller members in a larger ecosystem, ACT's members could be pressured to accept burdensome requirements to continue providing software to prime contractors. Additionally, ACT's members who develop specialized, domain-specific LLMs on top of existing, commercial platforms lack the economic power to encourage commercial platforms to adopt the clause's requirements.

The GSA should amend section (c) to only require prime contractors to take reasonable and practicable steps to assess and monitor subcontractors' competence to provide secure and reliable services. This change will increase competition and strengthen American AI innovation by encouraging more LLM developers to compete for government contracts and by reducing pressure on small businesses to accept non-standard commercial terms.

3. *The proposed clause imposes disproportionate burdens on small, micro, and non-technology businesses and should be paired with regulatory coherence and technical assistance.*

Because small businesses often lack the robust legal and compliance departments the proposed clause presupposes, its obligations will fall hardest on the smallest firms. A typical ACT member is a micro-enterprise with a team of fewer than 10 developers that has no dedicated compliance staff to monitor the countless automated changes that section (i) and the "material change" definition would require it to track, no in-house counsel to interpret the flowdown obligations in section (c), and little leverage to negotiate deviations from a prime contractor's terms. For ACT members, even a single outside-counsel review of the flowdown clauses or a single assessment of whether an open-weight dependency triggers the clause's foreign-control language can consume a meaningful share of their margin on a government contract. ACT notes that the government's own experience with compliance regimes provides insights in these costs:

- The Department of Defense estimated that a small entity would spend roughly $6,000 per year merely to self-assess and affirm the lowest level of the Cybersecurity Maturity Model Certification (CMMC) program, and more than $100,000 over three years to obtain a third-party certification at CMMC Level 2.[29]
- Industry estimates place the cost of a FedRAMP authorization at between $250,000 and several million dollars, with annual maintenance often exceeding $50,000.[30]

[29] See Cybersecurity Maturity Model Certification (CMMC) Program, 89 Fed. Reg. 83,092 (Oct. 15, 2024) (to be codified at 32 C.F.R. pt. 170). The Department of Defense's regulatory analysis estimated, for example, that a small entity would incur roughly $5,977 per year to self-assess and affirm compliance at CMMC Level 1, and approximately $104,670 over three years to obtain a third-party (C3PAO) assessment at CMMC Level 2.
[30] The FedRAMP program does not publish official cost figures, but industry estimates place total authorization costs at roughly $250,000 for a low-impact system to several million dollars for higher

Layering the proposed clause's undisclosed evaluation methodologies, decommissioning-cost exposure, and continuous change-reporting obligations on top of these existing regimes will price many small and micro businesses out of the federal LLM market entirely.

ACT notes that these compliance burdens are not limited to companies that build or host LLMs. Because the clause reaches any contractor whose offering processes government data through an LLM, its applicability can shift mid-performance when a firm adds an LLM feature to an existing product, drawing in professional-services firms, specialized consultants, and vertical software-as-a-service providers as prime contractors or subcontractors as LLM functionality becomes a routine component of ordinary software. ACT urges the GSA to account for these downstream small businesses and to calibrate obligations to the contractor's actual role and to the volume and sensitivity of the government data at issue.

Taken together, the features of GSA's proposal that ACT has identified (opaque evaluation methods, decommissioning-cost risk, absolute foreign-component language, and broad flowdowns) are poised to operate as barriers to entry that systematically favor a handful of large, well-capitalized incumbents that can absorb undefined compliance risk, staff continuous monitoring, and indemnify prime contractors against flowdown liability. Ultimately, GSA will be shrinking the pool of viable vendors by effectively excluding specialized small businesses best positioned to build performant, domain-specific models, leading to higher prices for the government and slower innovation. ACT strongly urges GSA to preserve a competitive base of small and mid-sized suppliers to keep costs competitive and as a matter of national security.

To help small businesses meet compliance burdens, we urge the government to pair any final clause with comparable technical assistance (plain-language implementation guidance, model flowdown language that small businesses can adopt without intense legal review, and a designated point of contact for small-business compliance questions). Such assistance would materially reduce the clause's disproportionate impact on smaller firms at modest cost to the government. ACT notes that the Small Business Administration and the Department of Defense fund APEX Accelerators (formerly Procurement Technical Assistance Centers) may provide a useful model for this focused effort.[31]

Finally, ACT notes that federal contractors that deploy LLMs already navigate a growing patchwork of AI requirements, including OMB Memoranda M-25-22 on AI acquisition and M-26-04 establishing the Unbiased AI Principles, the NIST AI Risk Management Framework,

---

baselines, with annual continuous-monitoring costs frequently exceeding $50,000. See, e.g., Vanta, How Much Does FedRAMP Authorization Cost?, https://www.vanta.com/collection/fedramp/fedramp-cost.

[31] See Procurement Technical Assistance Cooperative Agreement Program, 10 U.S.C. §§ 4951–4956; U.S. Small Business Administration, APEX Accelerators, https://www.sba.gov/local-assistance/federal-contracting-assistance.

and evolving Department of Defense and agency-specific guidance.[32] For small businesses, each additional requirement multiplies compliance complexity and legal uncertainty. ACT urges the GSA to harmonize this clause with existing federal AI guidance (including by expressly cross-referencing the NIST AI Risk Management Framework's risk-based approach) so that small businesses can operate under one coherent set of expectations in support of government goals and operations.

Relatedly, although GSA has advanced this clause through a Schedule refresh rather than ordinary notice-and-comment rulemaking, ACT urges GSA to undertake, before finalizing, the kind of small-entity impact analysis that the Regulatory Flexibility Act contemplates for significant rules.[33]

[32] See Office of Mgmt. & Budget, M-25-22, Driving Efficient Acquisition of Artificial Intelligence in Government (Apr. 3, 2025); Office of Mgmt. & Budget, M-26-04, Increasing Public Trust in Artificial Intelligence Through Unbiased AI Principles (Dec. 2025); Nat'l Inst. of Standards & Tech., AI 100-1, Artificial Intelligence Risk Management Framework (Jan. 2023).

[33] Cf. Regulatory Flexibility Act, 5 U.S.C. §§ 601–612 (directing agencies to assess the impact of rules on small entities and to consider less burdensome regulatory alternatives).

IV. Conclusion

As the GSA develops procedures for the government acquisition of LLM technologies, ACT encourages the GSA to disclose the criteria and methods contractors' LLMs will be evaluated against. As described in NIST's Risk Management Framework, we also encourage the GSA to enable contractors to use a risk-based approach to assess open-source components, globally operated systems, and change notifications to allow contractors to use common, commercial software solutions. These changes will encourage American innovation, foster competition, and will enable contractors to supply cost-effective and performant LLMs to the government agencies.

ACT appreciates the opportunity to submit these comments. We stand ready to work with the GSA and other stakeholders to develop procurement policies that strike the appropriate balance of protecting government data and encouraging American AI innovation.

Sincerely,

Brian Scarpelli
Senior Global Policy Counsel

Chapin Gregor
Senior Policy Counsel

Tess Huelskamp
Legal Intern

Association for Competitive Technology (ACT)
1401 K St NW (Ste 501)
Washington, DC 20005