Business Software Alliance Comments on the General Services Administration's Request for Comment on the Draft General Services Acquisition Regulation Clause on Large Language Model Artificial Intelligence Systems

July 23, 2026

The Business Software Alliance (BSA) appreciates the opportunity to provide comments on the General Services Administration's (GSA) Request for Comment on the Draft General Services Acquisition Regulation (GSAR) clause on Large Language Model (LLM) Artificial Intelligence Systems. BSA is the global trade association of the enterprise software industry, representing companies that are leaders in AI, cybersecurity, cloud computing, and other cutting-edge technologies.[1] We work in over 20 markets in the United States, Europe, and Asia, advocating for policies that build trust in technology so that every industry sector and the public can benefit from innovation. As a result, BSA members have unique insights into AI's tremendous potential to spur digital transformation and the policies that can best support the responsible use and adoption of AI.

BSA's core goal here aligns with the Trump Administration's goals: to accelerate the federal government's widespread adoption of AI. Our comments below identify the unintended consequences of the draft clause, which will thwart this important objective and limit federal access to cutting-edge AI products and services. BSA appreciates GSA's consideration of these important issues and the steps it has taken to better enable both US and non-US based companies to provide AI products and services to the government. However, BSA has significant continued concerns about the draft GSAR clause, including the need to limit its applicability to LLM service providers or operators unless they are developing and/or operating the system on behalf of the government, failure to recognize contractors' inability to "flow up" requirements to LLM developers, allocation of data ownership that significantly impedes contractors' intellectual property (IP) rights, and infeasible implementation of the unbiased AI principles. BSA articulated several of these concerns in its comments on GSA's initial draft of the proposed GSAR clause.[2] We highlight below recommendations to further enhance the draft GSAR clause, including:

- Clarify and revise the scope of the draft GSAR clause to limit it to LLM developers, expand the commercial product exception, and exclude use of retrieval mechanisms and knowledge bases. This clause should not apply to LLM service providers unless they design, build, deploy, and operate custom (i.e., non-commercial) LLM systems on behalf of the federal government because they will unlikely be able to comply with the clause;
- Remove the foreign-based restrictions on the use of open source components;
- Revise the GSAR clause's allocation of IP ownership and licensing rights, along with the data-related definitions, to align them with standard commercial licensing practices and avoid impeding contractors' IP rights;

---

[1] BSA's members include: Adobe, Alteryx, Amadeus, Asana, Atlassian, Autodesk, Avalara, Bentley Systems, Box, Cisco, Cohere, Cohesity, Dassault Systemes, Databricks, Datadog, Docusign, Dropbox, Elastic, EY, Graphisoft, HubSpot, IBM, Kyndryl, MathWorks, Microsoft, Notion, Okta, OpenAI, Oracle, PagerDuty, Palo Alto Networks, PTC, Rubrik, Salesforce, SAP, ServiceNow, Shopify Inc., Siemens Industry Software Inc., TrendAI, TriNet, Veeam, Workday, Zendesk, and Zoom Communications Inc.

[2] *See* Business Software Alliance Comments on the General Services Administration's Proposed Clauses for Solicitations and Contracts for Artificial Intelligence Capabilities, Apr. 1, 2026, https://www.bsa.org/policy-filings/us-bsa-comments-general-services-administration-solicitations-contracts-ai-capabilities.

- Revise flowdown requirements and exceptions;
- Improve implementation of the Unbiased AI Principles by limiting requirements to commercially reasonable efforts, recognizing shared responsibilities of LLM developers and government customers, and using widely available evaluation methods in lieu of undisclosed, bespoke government standards;
- Eliminate the regulatory-based modification disclosure requirement and, if retained, limit it to AI systems that have been specifically modified or uniquely configured at the request of a non-US government in a manner that materially deviates from its standard commercial configuration;
- Revise the compliance and reporting, change notification, and data portability requirements;
- Require contracting officers to tailor the application and requirements of this clause to the specific contract holder in accordance with the related Office and Management and Budget (OMB) guidance and EOs (e.g., allow contracting officers to not require contractors to comply with requirements that should not apply to the contractor). This will help ensure federal agencies do not lose access to existing AI and to ensure the President's AI-forward agenda and strong preference for commercial software is not thwarted by restrictive regulation; and
- Align explicitly with existing authorizations under FedRAMP/IL-5 frameworks to avoid duplication of effort and conflicting standards and enable contractor flexibility to meet changing needs.

I. Clarify Scope of Covered AI Technologies

The GSAR clause's application to "LLM AI systems" conflates different stages of the technology development and creates confusion about which products are covered. For example, it is unclear whether the GSAR clause applies to LLMs – the original models, AI systems leveraging embedded LLMs, data platforms offering access to LLMs, or a mix of all three. This is a very important distinction that lies at the core of determining whether a host of companies and software products are covered (and whether compliance is commercially possible). The GSAR clause appears to be intended to capture AI systems leveraging embedded LLMs or data platforms offering access to LLMs, but most requirements, except for data portability, relate to measures controlled by the original LLM developers and, therefore, do not align with the role of the potentially covered company in scope and their incapacity to comply with most of the requirements. GSA should clarify the scope, limit it to LLMs (but, as discussed in Section V below, in a flexible way that promotes disclosure and customer responsibility over rule-based compliance that will limit access to LLMs), and avoid imposing obligations on LLM service providers and operators that they have no ability to implement.

Even if limited to LLMs, the proposed scope would affect a wide range of commercial data and search platforms that are not themselves generative AI systems and that may not even be offered in connection with AI (e.g., the environment in which the model is deployed, operated, or monitored, as well as knowledge bases, APIs, monitoring sytems, and user interfaces). GSA should carve out non-AI functionality and explicitly exclude two scenarios: (1) non-generative AI models derived from generative AI models and (2) non-generative AI search and retrieval infrastructure that feeds a generative AI model. The proposed defintion expressly includes "retrieval mechanisms" and "knowledge bases," which could sweep in general-purpose, non-generative AI search infrastructure simply because a generative model leverages that information.

Importantly, the draft GSAR clause includes a common commercial product exception, but as BSA previously highlighted, the scope of that exception, which is based on the Advancing American AI Act – a law that pre-dates widespread availability of generative AI models, is not sufficiently broad to exclude a range of widely used commercial products. OMB's M-25-22 memo on government use of AI identifies a broader exclusion of widely available commercial products that GSA should leverage and further expand. GSA should also anchor the commercial exception in FAR 2.101, which sets forth the definitions of both "commercial products" and "commercial services," rather than introducing the

undefined and potentially ambiguous concept of "common commercial products." Notably, explicitly identifying software and cloud services as examples that fall within the scope of the exception could address the unnecessary inclusion of a wide range of commercial offerings. This approach would better align with the statutory preference for commercial products and commercial services under 41 U.S.C. § 3307.

In addition, we note that GSA has attempted to broaden the exemption from the GSAR clause's applicability by adding Section 539.71(b)(2), which provides an exemption when the "LLM functionality is incidental to the primary purpose of the core requirement being procured." We appreciate GSA's efforts to address the limited scope of this exception. However, the proposed language is tied to the purpose of the procurement rather than the nature of the product or service being procured. Importantly, this approach uses the wrong trigger for the assessment and would create an unpredictable case-by-case assessment of a purpose, which itself may be multifaceted and evolving. For example, a work management platform that embeds AI features for task prioritization or workflow automation can disable those features without affecting core contract performance. The "purpose of procurement" trigger does not reliably capture that distinction; the "nature of the product" standard does. As a result, we recommend that GSA revise the exception to include an array of circumstances that should fall within the exemption of the GSAR clause.

<u>Recommendation</u>: GSA should revise the scope of covered products by:

- **Limiting application to LLM developers, not entities integrating LLMs into other software systems. This should not broadly apply to a LLM system.**
- **Revising Section (a)(1)(ii)(A) to state: "The LLM is embedded in a commercial product or service (as defined in FAR 2.101), such as a map navigation system, software application, data or analytics service, or cloud service."**
- **Adding: "An LLM does not include non-generative machine learning or data-processing functionality – including non-generative models derived from or based on a generative model – or the general-purpose data infrastructure on which that functionality operates, where that functionality operates independently of real-time generative inference, even if used in conjunction with a generative model."**
- **Revise definition of "*Large Language Model Artificial Intelligence System"* to just "Artificial Intelligence Large Language Model" and eliminate reference to "integrated technical and operational environment in which the model is configured, deployed, operated, monitored, or made available for the processing of Government Data, including the model, hosting and access infrastructure, system prompts, configurations, knowledge bases, retrieval mechanisms, Application Programming Interfaces (APIs), user interfaces, guardrails, monitoring systems, and associated workflows."**

II. Remove the Foreign-Based Restrictions on the Use of Open Source Components

We appreciate that GSA has substantially revised the draft GSAR clause to enable more flexible use of AI products and services offered by non-US companies, US companies with global operations, and companies leveraging foreign-developed and/or open source components. We note, however, that the GSAR clause's foreign control restrictions on the use of open source continue to impose infeasible requirements and effectively impede the use of open source, which significantly impairs contractors' ability to provide innovative AI products and services to the government.

Section (f)(2)(iii) explicitly authorizes the use of foreign-developed components, including open-source components, and globally operated infrastructure dependencies so long as they do not violate the foreign control criteria in Section (f)((2)(i) and (f)(2)(ii), which specify the preference for US companies and restrict control by foreign governments. As BSA explained in Section II.B. of its comments on the GSA's initial draft clause, these foreign-based restrictions on open source effectively ban the use of

open source software because of how it is compiled and the resulting practical infeasibility of tracing the origin of all components or assessing the degree of control exercised by a foreign entity.

Definitively tracing the geographic origin of an LLM requires a level of detail that is not typically provided for commercial products and commercial services and may not be readily available under the commercial terms under which an LLM is licensed by its developer. Indeed, in many instances, contractors will be unable to definitively trace the geographic origin of open source model contributors, making compliance technically and administratively difficult at best and, in some instances, infeasible. Indeed, this overbroad restriction fails to account for the "hidden lineage" inherent in the modern, composable AI ecosystem. Routine development practices—such as fine-tuning open source models, distilling them into smaller, faster versions, and iterative re-labeling—are standard and legitimate industry practices that systematically obscure a model's original lineage. While these connections persist technically within model weights, they are frequently absent from standard documentation and dependency manifests. Thus, the foreign-based restrictions ignore the reality that models are often built upon a complex, multi-layered web of open-source contributions. Imposing this restriction would force companies to perform deep-layer forensic analysis on model weights that, in many cases, for the reasons stated above, may not yield definitive origin data. This creates a currently impossible compliance standard that does not reflect the dynamic, global nature of the AI supply chain and will ultimately discourage the use of innovative, open-source-derived technologies in government contracts. Notably, even in circumstances where contractors have not sought to include open source components, the proposed requirement would impose substantial operational burdens on contractors to trace and certify the national origin of all components within complex AI systems, which is not technically feasible given the widespread use of open-source components and globally distributed development processes.

Further, the clause also suffers from several other defects. First, it fails to provide an objective standard for determining when a model is deemed to derive from a foreign-origin component. Modern development techniques, including fine-tuning, distillation, synthetic-data generation, and other derivative training methods, create varying degrees of impact from upstream models, but the clause does not appear to define the endpoint of lineage analysis. As a result, the proposed approach lacks a knowable compliance standard, resulting in contractors, contracting officers, and auditors reaching different conclusions and leaving contractors with insufficient information about how to demonstrate compliance with a standard that is impossible to meet. Second, there is no explicit carve-out for internal research, evaluation, or benchmarking activities conducted in isolated environments with no connection to production systems or government data. Contractors routinely evaluate a broad range of AI models, including non-US models, as part of standard R&D practices such as competitive benchmarking, safety testing, and model selection. These activities are entirely separate from contract performance and serve legitimate purposes, including ensuring contractors can identify and deploy the best American AI systems for government use. Without clarification, contractors face the untenable choice of either halting critical research practices, which are necessary to remain competitive, to avoid compliance risk or choosing not to sell their products on MAS. Third, the reference in Section (f)(2)(iii) to "incidental" foreign-developed components threatens to further limit their use, as it may raise questions about whether the integration of foreign models is "incidental" or core to the software functionality. Accordingly, we recommend deleting this term. Finally, GSA should ensure that the clause's reference to control or compulsion by a foreign government is applied in a manner that is consistent with existing statutory approaches and aligns with existing frameworks.

In short, the proposed clauses make a US company's AI system that integrates open source software or an open weight model in its initial development but otherwise develops the AI system in the United States and hosts the government's data in a compliant, controlled environment, ineligible for government purchase. Such an unacceptable result would run afoul of the Trump Administration's support of the open source AI ecosystem, which levels the competitive playing field, enables companies to provide AI services at lower cost, and allows companies to customize AI systems to deliver more innovative, tangible benefits. The inability to leverage open source software or open model

weights would also contravene the Administration's broader IT modernization objectives and limit the US government's ability to purchase low-cost, state-of-the-art commercial AI solutions.

Recommendation: Remove the foreign-based restrictions on open source components. The clause should expressly permit open source and open weight models, including foreign-origin models, so long as applicable security and data protection safeguards are met. The FAR and FedRAMP comprehensively address these security and data protection requirements already, and further restrictions do not advance government access to cutting-edge AI or data security/protection.

III. Revise Data Definitions and Clarify Contractors' IP Rights to Align Them With Standard Commercial Licensing Practices

We urge GSA to revise the GSAR clause to align the allocation of IP and ownership rights, including the data-related defintions, with standard commercial IP regimes and norms. We recognize that the government has a legitimate interest in maintaining ownership rights over its own data. However, we have significant concerns about the current approach reflected in the GSAR clause because it significantly impedes' contractors' IP rights by departing from standard commercial practices, including outlining overbroad data definitions, and capturing ordinary commercial product development rather than government-funded bespoke work. To better align with commercial practices and the statutory preference for commercial IT procurements, GSA should allocate IP and ownership rights based on the adoption of commercial terms and conditions, supplemented only with contract clauses that are either necessary to implement applicable legal rights to commercial products or services, or consistent with standard commercial practice (see 41 U.S.C. § 3307(e)(2)).

Consequently, as set forth below, we urge GSA to:

(1) Remove the overbroad assertion of government ownership over custom development, which ignores the extent to which those developments are developed for and implemented across a range of commercial services;
(2) Expand the scope of contractors' rights to explicitly protect their proprietary code, documentation, knowledge base articles, and ownership of data outputs, including service-generated and derivative data; for commercial vendors, allow them to license according to their own commercial terms, consistent with FAR Part 12;
(3) Enable use of government feedback and remove the assertion of government ownership over automated feedback;
(4) Revise the approach to background data to recognize contractors' IP protection after changes have occurred beyond the original form;
(5) Enable the use of government-related data such as usage data and model inputs and outputs, to develop, improve, operate, and support a contractor's products or services and for security, abuse prevention, reliability, troubleshooting, maintenance, and service administration purposes;
(6) Revise the definition of government data to expressly exclude de-identified or aggregated operational metadata, telemetry, and system diagnostics that contain no government-identifiable content, consistent with NIST de-identification standards;
(7) Delete the reference to an "irrevocable" license for the government, which conflicts with standard contracting practices;
(8) Revise the post-termination retention obligations so they do not sweep in protected IP or impose other operational challenges;
(9) Expand the rights to use or process government data to prevent unintended operational limitations; and
(10) Delete the provision referencing the transfer of third parties' IP rights, which contractors are unable to control or convey.

A. Custom Development

First, the definition of "custom developments" should be narrowed to apply to deliverables developed exclusively for and funded by the government. As drafted, it captures "modifications, customizations, configurations, or enhancements" that may reflect a provider's ordinary commercial development, including model fine-tuning that improves the underlying commercial offering. In standard commercial practices, the contractor generally owns the background IP and custom developments, while the government retains ownership of its own data, and obtains a license to use custom developments, including configurations, enhancements, work product and deliverables, during the duration of the contract term.

The assertion of the government's ownership of "custom development," which is so broad that it could capture software tools, workflows, routine configuration activities, or customer-specific utilities created for government customers, does not recognize the extent to which these developments are based on contractors' IP and broadly integrated across a range of commercial services. This novel proposal conflicts with established federal data and software rights frameworks, which could create confusion in how agencies apply those regimes across their portfolios. Such an approach will disincentivize participation in the federal marketplace because it undermines companies' private sector investments. It is imperative that platform-wide updates to contractors' products and services are not misclassified as government-owned assets, restricting contractors' ability to leverage those improvements for other customers. For example, if a contractor enhances a product for fraud detection across a global workflow, with minor implementations for a government customer, it should not be restricted from leveraging the fraud detection feature in its platform in other commercial settings.

In addition, the assertion of rights over custom development creates uncertainty about the line where government ownership ends and typical proprietary rights begin, which could require contractors to create bespoke solutions for the government that are not updated as frequently simply to preserve appropriate ownership rights in commercial contexts, undermining the aim of modernizing federal IT with more cutting-edge commercial solutions. To the extent the government seeks ownership of custom development, it should refer to the bespoke work product specifically funded under the federal contract, not platform-wide features or capabilities deployed to multiple customers.

B. Background Data

The issues highlighted above also implicate the GSAR clause's definition of background data, which includes, among other things, *pre-existing* proprietary content, reference materials, and knowledge bases "referenced, retrieved, augmented, or otherwise incorporated into the LM's processing or outputs through any enrichment mechanisms." We appreciate GSA's addition of background data as a defined term. However, the reference to pre-existing IP, coupled with Section (e)(1)(v)'s statement that the contractor retains ownership of background data "in its original form," unduly restricts contractors' IP rights in this data, as the underlying base model may change over the course of the contract, and that category of information continues to warrant protection.

C. Feedback

Section (e)(1)(ix) of the GSAR clause states that the government owns feedback that it provides to the contractor with respect to the LLM or custom developments, regardless of whether the feedback is generated through automated processes, and prohibits the use of that feedback for system improvement or any other purposes where feedback includes government data or confidential information. However, feedback rarely benefits only one customer, and contractors should be able to freely use it for AI system improvements, consistent with standard commercial practices. Further, the use of automated performance signals is a routine part of product operation and should not transfer to the government. Usage data/telemetry that is included in internal logs must be retained in contractor usage logs for FedRAMP compliance, and it may also be intermixed in multi-tenant logs that cannot readily be deleted upon demand. Consequently, any assignment of government ownership for feedback should exclude automated operational or performance telemetry, including where government data is incidentally retained. Further, GSA should clarify that any feedback provided by the government does not grant the government ownership rights over underlying technologies, derivative works, or product improvements developed by the contractor in response to the feedback, as such rights could inhibit ongoing innovation and commercial product development. Conversely, GSA should note that feedback does not – and should not – include the right to use customer confidential information.

D. Data Outputs and Inputs

In addition, the overly broad definition of data outputs, which has not changed significantly since the initial draft, does not sufficiently address rights in such data or align with standard commercial practices. GSA should revise the definition of data outputs and data inputs so that it does not interfere with contractors' IP rights.

The GSAR clause's defintion of data output includes "responses, results, analyses, anonymized data, derivative data, metadata, logs, synthetic data, and any other output or action produced by the LLM, regardless of whether such output incorporates or is derived from background data," exempting "technical system-level data that contains no government information or government usage context, such as performance metrics, token counts, and processing times." The definition is overbroad. The definition of data output should be limited to the outputs that the government generates from its use of the LLM. The government should allow the contractor to collect and use usage data and LLM inputs and outputs to develop and improve the contractor's services under certain circumstances discussed below.

First, the data output definition fails to take into account the likelihood that third-party data, which may not be owned by the government or the contractor, could be included in data inputs. More generally, GSA should clarify that the government's ownership of data outputs only extends to data outputs that do not already belong to a third party. This distinguishes between IP existing before and after the output was created and specifies that companies are only conveying to the government ownership in new IP contained in a data output, not someone else's pre-existing IP that happens to be contained in the data output. This distinction respects existing IP rights and supports continued innovation.

Second, the definition includes system-generated information, such as metadata, logs, telemetry, and system actions, which are typically owned by the system provider. The inclusion of data derived from protected background data further expands government ownership in a way that substantially impedes contractors' IP rights. Notably, the inclusion of any "action produced by the LLM" in the data output defintion creates several challenges. Many AI products generate recommended actions from a finite, predefined set of permissible actions within the product's functionality. System options, such as "view status update" for a particular project, are not creative or unique outputs generated from government data but rather system commands. Under the current language, the government would obtain ownership rights not only in the outputs generated in response to government prompts, but also in the

underlying action taxonomy in the system commands themselves. Such an assertion would be overbroad and would effectively transfer ownership of functional components of contractors' commercial products. In addition, requiring contractors to maintain separate product configurations for government and commercial customers to avoid this outcome would impose significant engineering and compliance costs that would chill participation in federal AI procurement. The assertion of government ownership over operational metadata and telemetry also poses additional security-related concerns, as contractors leverage this information to secure their products. Further, if technical telemetry associated with a federal environment is treated as government data, the clause could trigger data localization, human access, deletion, portability, and flowdown obligations even where the telemetry contains no government content or identifying information. Applying those requirements to routine security, reliability, capacity-management, and service-improvement data could require substantial re-architecture of shared SaaS infrastructure and impair providers' ability to detect security incidents, maintain service reliability, and support customers.

Third, the exemption from the data output definition should also exclude other vendor proprietary information, including reference information or documentation. For example, AI products may assist users in troubleshooting by summarizing the contractor's documentation or leveraging retrieval augmented generation (RAG) to provide a reliable information source to respond to the prompt. However, the contractor should retain ownership in those proprietary documentation and reference sources that may appear within or be cited by an output. Similarly, GSA should revise the definition of data input to exclude vendor-specific information provided to the government, such as product documentation and user guides, as they are critical for reliability and security monitoring.

In sum, we recommend that GSA revise the definition of data output and input to exclude third-party data, exclude system actions, derivative data, or other service-generated metadata and logs; clarify that aggregation, de-identification, and security operations do not, by themselves, convert otherwise non-government telemetry into government data; delete the reference to data derived from contractors' protected background data; and expand the exemption to exclude other vendor proprietary information, such as reference information or documentation, or any other background IP or information existing prior to entry into the contract or developed independently by the contractor.

E. Product Improvement

Section (e)(3)(i) of the GSAR clause, which is substantially similar to the initial proposed clause, outlines the prohibited uses of government data, which includes "training, fine-tuning, or otherwise improving an LLM, including those operated by third parties, or to develop or improve the LLM(s) for any other customers or any commercial or noncommercial purposes." We appreciate the aim of appropriately protecting government data that is accessed and used in connection with the performance of the contract. However, we note that the proposed clause may unnecessarily limit the use of data for improvement of the product under the contract, as well as for commercial settings where contractors have implemented appropriate safeguards, such as adequate security controls and robust de-identification mechanisms, to protect the confidentiality of government data. Further, certain proposed requirements, such as the segregation or removal of government data from AI systems, do not reflect the technical realities of modern machine learning systems. Once data has been incorporated into model training, it is not technically feasible to isolate or remove specific data elements from trained models.

Moreover, the proposed approach will be difficult to operationalize because of the breadth of the data definitions and the various ways data might be integrated into the AI system. For example, as discussed above, "data outputs" is defined broadly to include all data generated by an AI system, including metadata, logs, derivative data, and other types of data. While system-level data containing no government data or government usage context is carved out of the definition, many contractors cannot control the extent to which users input their own data into fields that may ultimately end up in

system-level or operational metrics. Typically, vendors ask customers to avoid inserting personal, confidential, or sensitive information in fields that may be captured in logs and other metadata for that reason. In addition, "government usage context" is a vague term that could broadly apply to almost any metrics related to the government's use of an AI system. GSA's restrictions on the use and management of government data, which includes data outputs, including with respect to use for operational purposes, will create significant technical and compliance difficulties for contractors and impede visibility for monitoring purposes.

Further, even if GSA retains the broad restriction on product improvement for commercial purposes, it should revise the language to enable improvement of the product or service provided under the government's contract and of the category of products used by the government and commercial users alike. Currently, the language "or any commercial or non-commercial purposes" would restrict a contractor's ability to use the government data to train, fine-tune, or otherwise improve or customize a model for the government under the performance contract. For example, in the generative AI context, contractors may improve the product through prompt engineering or testing and evaluating retrieval augmented generation (RAG) efforts to prevent drift or hallucinations, which enhance accuracy and performance. In the nongenerative AI context, this restriction has no purpose because there is no real risk of disclosure of government data, such as an image recognition model supporting computer vision to identify objects. Presumably, GSA does not intend to limit the US government's ability to obtain the full suite of benefits of the AI product or service it is purchasing, as enabling the use of government data to improve the product or service it is using will be an important way to maximize its investment.

F. "Irrevocable" License

Section (e)(2) provides the government with an "irrevocable" license for the duration of the work in the contract or task/delivery order. However, this is inconsistent with standard commercial licensing practices, foundational contract law, and established government contracting principles, which typically provide for termination rights under defined conditions. Moreover, the term "irrevocable" may be interpreted as limiting a contractor's ability to suspend or terminate access in circumstances involving customer breach, misuse, or other situations where service termination would otherwise be appropriate. Further, a contractor that is licensing a third-party LLM from an upstream supplier can't guarantee that its license from the supplier will never be terminated, as discussed more fully in Section IV below. As a result, GSA should remove the "irrevocable" requirement and instead align licensing terms with standard commercial frameworks.

G. Data Deletion Requirements

The data deletion requirements, including in Section (e)(3)(iii) and (e)(vii), are overbroad and should be more narrowly scoped. For example, they fail to account for the need to retain data for legal, compliance, audit, security, backup, or litigation-hold purposes. Notably, the deletion requirements do not capture how inputs may be reasonably retained in contractor usage logs and product telemetry in commercial systems (even within the FedRAMP boundary) and that are routinely retained for legitimate business purposes, including FedRAMP retention requirements. Further, the requirements do not appear to accommodate standard data deletion schedules. Consequently, we recommend that GSA revise the deletion requirements to account for the broader impact on commercial systems, particularly with respect to use logs, and avoid creating conflicts with other legal requirements.

H. Expand Other Rights to Use or Process Government Data to Prevent Unintended Consequences

The GSAR clause's IP section also includes restrictions on use and processing of government data that may create unintended operational limitations. For example, Section (e)(1)(iii) enumerates certain permitted uses, but it does not expressly authorize use of government data for security purposes, which

is important to ensure appropriate cybersecurity measures. In addition, Section (e)(1)(ii) restricts the entities that may process government data, but it does not acknowledge that subprocessors or service providers also routinely process data in connection with providing the contracted services, even where they are not performing the specifically defined LLM-related roles. As a result, GSA should revise these provisions to enable the use of government data for security purposes and the processing of government data by subprocessors and service providers.

I. Logging, Traceability, and Government Access

Sections (e)(4)(ii) and (e)(4)(iv) require audit logging that avoids capturing or displaying actual government data, tools enabling detailed government records of processing activities, and summaries of intermediate processing actions, model-routing rationale, and retrieval methods. These requirements should enable meaningful government auditability of the contracted use case, rather than effectively requiring a generalized transparency or explainability program developed for other regulatory regimes. Without that distinction, the clause could be read to require disclosure of proprietary reasoning, model mechanics, routing logic, or non-Government service information that is not necessary for the government to investigate, govern, or troubleshoot its use of the service. It also could require new logging architectures that are not commercially available or technically feasible and that could create additional security, privacy, or data-exposure risks.

GSA should clarify that logging and traceability obligations require audit records sufficient for the Government to investigate, govern, secure, and troubleshoot its contracted use case. The clause should not require generalized model explainability, chain-of-thought, proprietary internal reasoning, disclosure of internal model mechanics, or new logging architectures beyond commercially available and technically feasible audit capabilities. Government-accessible logs should be logically segregated and configured so that government Data remains accessible only to authorized government users, except as necessary for authorized support, security, or incident response.

Recommendations:

- **Section (b): Revise the definition of custom developments to exclude "any configuration, customization, enhancement, or model improvement that the Contractor makes generally available, or could make generally available, as part of its commercial product or service."**
- **Sections (b) and (e)(1)-(3): Revise the definition of data outputs to be limited to outputs that the government generates from its use of the LLM and to exclude data owned by third parties and system-generated data, including usage data, metadata and logs. Revise the IP Rights section so the government allows the contractor to collect and use usage data (e.g., data and telemetry collected by the contractor related to the government's use of the services) to develop, improve, operate, and support their product and services. In addition, revise the IP Rights section so the government allows the contractor to collect and use LLM inputs and outputs to develop and improve the contractor's services, assuming such use (1) is only for the benefit of the government or (2) does not identify the government or disclose the government's confidential information to third parties.**
- **Section (e)(1)(ix): Delete Section (e)(1)(ix), authorize contractors' and service providers' non-exclusive, worldwide, royalty-free right and licenses to use feedback for AI system improvements, and ensure that any ownership transfers or restrictions on feedback do not include automatically generated system data.**
- **Sections (e)(4)(vii) and (e)(1)(i): Delete the reference to custom developments in Section (e)(4)(vii) and (e)(1)(i), consistent with the discussion above regarding the contractors' retention of IP rights.**
- **Section (e)(1)(v): Delete "in its original form" in Section (e)(1)(v) to appropriately protect contractors' rights in background data and other IP.**

- **Section (e)(3)(iii) and (e)(3)(vii): Revise Section (e) to exempt circumstances where data is retained for other legal, security, or resilience purposes or that may be commercially infeasible to implement.**
- **Section (e)(2): Delete the reference to an "irrevocable" license in Section (e)(2).**
- **Clarify that aggregation, de-identification, and security operations do not, by themselves, convert otherwise non-government telemetry into government data.**
- **Section (e)(4)(ii) and (e)(4)(iv): The clause should not require generalized model explainability, chain-of-thought, proprietary internal reasoning, disclosure of internal model mechanics, or new logging architectures beyond commercially available and technically feasible audit capabilities.**

IV. Revise Flowdown Requirements and Exceptions

BSA appreciates GSA's shift to a role-based flowdown structure, which is a meaningful improvement over an undifferentiated flowdown obligation applied to every subcontractor regardless of role. However, BSA members continue to have significant concerns about how the flowdown mechanism will operate in practice. We highlight four issues below and propose targeted fixes for each. The challenges associated with flowing down the requirements in the draft clause highlight why the clause itself should be limited to LLM developers.

A. Applicability Exceptions Should Extend to the Flowdown Obligation

As drafted, the applicability exceptions in Section (a)(1)(ii) determine whether the clause reaches the prime contractor, but the flowdown obligation in Section (a)(2) is a separate paragraph with no cross-reference to those exceptions. As a result, a subcontractor, service provider, or supplier whose role would independently qualify for an exception under Section (a)(1)(ii), for example a common commercial product exclusion or an incidental use exclusion, has no textual basis to claim that exception once flowdown is triggered against the prime. This creates an unintended result in which entities GSA itself intended to exclude from the clause could still be bound to it through flowdown. BSA urges GSA to revise Section (a)(2) to state: "Flowdown is not required or allowed to any subcontractor, service provider, or supplier whose role independently qualifies for an exception under paragraph (a)(1)(ii)." The carve-out should extend to suppliers specifically, since the four defined roles do not otherwise capture upstream suppliers that may independently qualify for an exception.

B. Flowdown – or In Practice, "Flow Ups" – Cannot Be Guaranteed Where the Contractor Lacks Contractual Leverage

As one of the clause's most significant defects, it requires the flowdown of requirements to LLM developers and other parties that the contractor cannot contractually bind. The clearest example is the LLM Developer role: contractors generally have no contractual leverage over LLM Developers, who are typically not a party to, and often have no relationship with, the underlying government contract. This concern is not limited to open source or open weight model publishers, although those cases are the clearest examples, since open source developers (to the extent they are active and contactable) will not agree to government specific flowdown terms under any circumstances. The same leverage problem can arise for any role in the flowdown structure, including the System Operator, System Integrator, and Service Provider roles, whenever the contractor lacks a direct contractual relationship with the party performing that role. BSA members have raised this as a defect in the entire flowdown framework: contractors cannot control what other companies agree to, and requiring them to guarantee flowdown compliance from parties they cannot bind is not workable regardless of scope or carve-out language. At minimum, a contractor that does not own or control the underlying foundational model should not inherit LLM developer obligations under any flowdown structure GSA retains.

Note, unlike subcontractors that are hired to fulfill tasks on a particular government contract and where compliance with those terms are understood to be prerequisites to contracting, contracts with LLM developers are not typically entered into on a contract-by-contract basis for LLM operators and providers, and the use of open-weight models falls completely outside of this subcontractor framework. This is why LLM operators and providers underscore that they cannot flow down these requirements to LLM developers in almost any instance.

Members of Congress have raised a related concern in a June 11, 2026 letter to GSA Administrator Ed Forst. Senators Roger Wicker, Todd Young, and Ted Budd stated that the clause "exposes prime contractors to liability for compliance of their AI subcontractors, including those with no direct contractual relationship with a federal department or agency," citing Section 1824 of the FY2026 National Defense Authorization Act, Public Law 119-60. Section 1824 added 10 U.S.C. § 3459, which prohibits the Secretary from requiring that a clause be included in a subcontract for the acquisition of commercial products or commercial services unless that clause appears in the Defense Federal Acquisition Regulation Supplement under 10 U.S.C. § 3452. That limitation binds the Pentagon, not GSA, and does not take effect until it publishes the required list. However, GSA's Federal Supply Schedules, GWACs, and OASIS+ vehicles are used extensively by Pentagon and Intelligence community components, which account for a substantial share of the orders placed under those vehicles. Once the Pentagon publishes its list, contracting officers ordering under GSA vehicles will be independently barred by 10 U.S.C. § 3459 from requiring flowdown of the LLM Developer clause, or any other role-based flowdown clause, to a subcontract for commercial products or commercial services unless that clause appears on the list. Given the scale of the Pentagon's reliance on GSA vehicles for AI-enabled commercial products and services, this is not a narrow edge case.

This creates a structural problem for GSA's clause that extends well beyond Pentagon orders alone and runs counter to the Administration's broader direction on commercial terms. A single GSA contract vehicle typically serves both the Pentagon and civilian agencies under the same underlying commercial terms. If mandatory flowdown becomes unenforceable for Pentagon orders while remaining nominally mandatory for civilian agency orders placed under the identical vehicle, contractors would be forced to maintain two incompatible compliance regimes for the same commercial relationship with the same LLM Developer, depending only on which agency happens to place the order. That outcome is not administrable and does not reflect how contractors structure their commercial supply chains. The Administration has directed the use of commercial terms through Executive Order 14271, Ensuring Commercial, Cost-Effective Solutions in Federal Contracts. The clause as drafted does not reflect that direction: it would effectively create two different products, one for GSA's use in civilian procurements and one that complies with the statutory limitations applicable to Pentagon procurements. A single GSA clause cannot function as intended if it becomes selectively unenforceable for the Pentagon's share of orders while remaining nominally binding for every other agency that orders under the same contract vehicle. In short, this could incentivize agencies to go around GSA to avoid the conflict between GSA's restrictive contract language and the Pentagon's statutory requirement to use commercial items and services.

C. Attestation and Due Diligence Standards Require Clarification

Section (d)(3)(ii) allows a contractor to demonstrate compliance through attestation from an individual with appropriate authority representing the LLM Developer, System Operator, System Integrator, or Service Provider. The clause does not specify who qualifies as an individual with appropriate authority, or what form that attestation should take. This is a meaningful gap in an otherwise workable compliance mechanism. BSA requests that GSA provide a standardized attestation template, similar to the representation format used in FAR 52.204-26(c), and clarify what constitutes appropriate authority to attest on behalf of each defined role.

We also have related concerns about the due diligence obligation in Section (d)(1). One potential solution is to add a qualifier such as “consistent with standard commercial practices” to avoid an overly vague requirement to exercise due diligence over the LLM Developer, System Operator, System Integrator, and Service Provider roles. Another issue is how companies will demonstrate compliance under Section (d)(3), since the clause does not specify a frequency or trigger for that obligation. Both of these concerns point to the same underlying problem: Section (d) imposes due diligence and compliance demonstration obligations without anchoring them to an objective, commercially recognized standard or a defined cadence. BSA recommends that GSA revise Section (d)(1) to state that the contractor must exercise due diligence “consistent with standard commercial practices” in selecting and overseeing the LLM Developer, System Operator, System Integrator, and Service Provider roles. GSA should also clarify the frequency or trigger for demonstrating compliance under Section (d)(3).

BSA further recommends that GSA explicitly enumerate the forms of documentation that constitute a valid attestation, including: FedRAMP authorization packages, SOC 2 Type II reports, trust-center materials, marketplace terms of service, provider compliance documentation, and written certifications from an authorized representative of the relevant entity. Additionally, GSA should codify that a contractor's reasonable reliance on such documentation satisfies its due diligence obligation under Section (d) absent actual knowledge of noncompliance. This "actual knowledge" standard is essential to prevent contractors from bearing strict liability for third-party noncompliance they cannot detect through commercially reasonable means and reflects the standard applicable in analogous supply-chain compliance frameworks.

D. The Closed List of Defined Roles May Improperly Exclude Other Subprocessors

Section (e)(1)(ii) states that only LLM Developers, LLM System Operators, LLM System Integrators, and LLM Service Providers may receive and process government data. A plain reading of this provision would exclude any subprocessor or service provider that processes government data while delivering the contracted service but does not independently perform one of those four LLM specific roles, for example a cloud hosting provider or a customer support subprocessor with no LLM specific function. That result is unlikely to reflect the drafters’ intent, since ordinary, non-LLM subprocessors routinely need to process government data to deliver commercial services generally, independent of any LLM functionality. BSA recommends that GSA clarify this provision to allow for any subprocessor or service provider to process government data.

Recommendations:

- **Revise Section (a)(2) to state: “Flowdown is not required or allowed to any subcontractor, service provider, or supplier whose role independently qualifies for an exception under paragraph (a)(1)(ii).”**
- **Eliminate the mandatory flowdown of the LLM Developer clause and any other role-based flowdown clause. Instead, the government should only require disclosure of the LLMs used directly by the customer for the services to allow the government to make an informed choice about whether to use the LLM or LLM system.**
- **Provide a commercially reasonable standardized attestation template and clarify what constitutes appropriate authority under Section (d)(3)(ii).**
- **Revise Section (d)(1) to require due diligence “consistent with standard commercial practices,” and clarify the frequency or trigger for demonstrating compliance under Section (d)(3).**
- **Clarify that Section (e)(1)(ii) is not intended to exclude ordinary, non-LLM subprocessors and service providers from processing government data and revise the provision accordingly.**

V. Improve Implementation of Unbiased AI Principles

GSA aims to implement the directives in OMB's M-26-04 memo on Increasing Public Trust in Artificial Intelligence Through Unbiased AI Principles, which requires federal agencies to ensure procurement contracts address compliance with the unbiased AI principles—"truth-seeking" and "ideological neutrality"—articulated in President Trump's Executive Order 14319. However, the draft GSAR clause (1) creates operational frameworks that are not feasible to implement, (2) fails to set forth consistent, objective, and predictable requirements for contractors, and (3) inappropriately assigns liability to contractors both in the False Claims context and for decommissioning costs for issues associated with this requirement. Further, the US government already has obligations addressing these issues, and any standards should be grounded in existing laws and policies that account for the US government's obligations with respect to the acquisition of commercial products and commercial services.

For example, Section (j)(1)(3) requires contractors to enhance detection and mitigation of trustworthiness or bias, but the ideological neutrality principle suggests that contractors may not implement measures that would filter AI system outputs. This contradiction creates an impractical operational obligation that contractors will be unable to meet. Further, compliance with the proposed clauses could result in reduced performance and lower accuracy, which could itself result in biased outcomes, subjecting the government to legal liability. Current technical realities do not enable guarantees of factual accuracy, as they are inherently probabilistic. Further, "nonpartisan" and "neutral" are not sufficiently clear guideposts that can be translated into infallible technical measures that prevent users from eliciting non-neutral responses. Accordingly, we recommend GSA change this standard to reflect a more feasible "commercially reasonable efforts" standard, though we underscore that if contractors don't know how to interpret the principles, they will still be unable to configure them to be compliant.

Notably, Section (j)(3) fails to recognize that LLM system providers have no capacity to change LLM performance, as those decisions lie with the LLM developer. It also fails to account for the fact that model selection and use are increasingly a customer responsibility or a shared responsibility, not exclusively the contractor's. Thus, GSA should revise the Section (j) to either require the contractor to implement commercially reasonable measures to address this issue or enable the government to do so and make explicit that compliance is achieved by a good faith effort to do one or both. Software documentation is commonly used to explain shared responsibilities in other contexts (e.g., data security in the cloud) and is a sufficient solution here. This is critical as it will enable the government to retain access to LLMs and obtain experience designing and using LLMs in a compliant manner.

In addition, Section (j)(2) requires contractors to provide automated tooling for the government to run automated assessments of the AI system to evaluate, among other things, bias, truthfulness, and ideological content, but many contractors of cutting-edge AI software platforms do not currently have the technical ability to provide this function. Further, the proposed clause ignores the OMB M-25-22 memo's direction to agencies to enable vendors to complete their own testing and evaluation where appropriate. Along the same lines, it is not clear whether contractors would be required to perform custom development to create this tooling, and how it should do so if not provided the evaluation details, as discussed below.

Importantly, Section (j)(2)(iii) also indicates that the government will use its own benchmarks to evaluate compliance with the unbiased AI principles, and that the government is not required to disclose or provide access to the underlying data, methodologies, or systems for its assessments, except where there is an adverse action. The draft clause does not include any limitations on scope, frequency, or methodology and could create operational disruption, risk IP exposure, and duplicate existing frameworks, such as FedRAMP. Such an approach creates significant uncertainty for contractors, who will be unable to build products according to known and objective specifications, or invest soundly in offering services to the government because of the unpredictability associated with bespoke, subjective, and invariably inconsistent agency evaluations of their products that prevent them from designing systems that can reliably meet compliance expectations. In lieu of this approach, GSA should enable contractors to leverage industry or

other widely accepted, objective benchmarks to help assess whether their development of AI systems will satisfy the unbiased AI principles. Where applicable, agencies could rely on standardized, publicly documented evaluation suites, consistent with existing protocols, including FedRAMP, rather than bespoke, undisclosed tests. The evaluations should be scheduled with reasonable advance notice, limited in frequency, and conducted in a manner that protects contractors' IP and trade secrets. This will foster trust, enable compliance, and reduce uncertainty for contractors.

Further, in light of the uncertainty associated with evaluation benchmarks that are not based on known, widely accepted technical standards, attributing decommissioning costs to contractors for AI systems the government deems non-compliant appears to be an unreasonable cost-shifting mechanism that creates a significant financial contingent liability that will discourage contractors from providing low-cost, innovative, and routinely updated AI solutions to the US government.

Recommendations: We recommend that GSA provide a more straightforward implementation of the OMB M-26-04 memo and simply advise contractors to use commercially reasonable efforts to implement (and/or allow government customers to implement) the unbiased AI principles. If the contractor does not develop the LLM, the clause should recognize the shared responsibility of compliance, and the government, not the contractor, should ultimately be responsible for complying with the unbiased AI principles with LLMs it operates. The clause should not require the contractor to guarantee the unbiased AI principles, as this is not possible for contractors. The clause should only recommend that contractors provide safeguards, disclosures, or documentation and perform testing, to the extent commercially reasonable and within the contractor's control.

In addition, we recommend that GSA omit bespoke government benchmark and automated assessment provisions, including the contractor-provided tooling, and instead enable contractors to leverage objective commercial or widely accepted benchmarks for responsible AI model development to demonstrate alignment with the unbiased AI principles. At a minimum, any government assessments should be conducted with technically feasible and secure existing protocols, including FedRAMP protocols, scheduled with reasonable advance notice, be limited in frequency, and conducted in a manner that protects contractors' IP and trade secrets. Further, we urge GSA to remove the contractor liability for decommissioning costs for AI systems deemed non-compliant given the variability associated with generative AI system outputs and the subjective and unpredictable nature of a non-compliant determination.

VI. Eliminate or Revise Reporting Requirement for Regulatory-Based Modifications to AI Systems

Section (f)(3) requires contractors to disclose whether the LLM has been modified or configured to comply with non-U.S. federal laws or policies. This disclosure requirement is overbroad and unworkable. Generally, globally competitive commercial AI products are designed to comply with US laws and other jurisdictions for maximum scalability across major markets and to support global customers, including US companies with global operations. This global reach enables US technology to compete and become embedded in as many markets as possible, cementing US leadership. Many responsible AI practices, including transparency mechanisms and human oversight controls, reflect industry best practice, customer requirements, and US legal obligations simultaneously. As a result, it would be impractical to disclose each time a design choice with overlapping aims is made, particularly where there is no adverse effect on compliance with the federal contract. Moreover, it creates a chilling effect on proactive investment that does not serve the government's interest in reliable AI.

GSA should eliminate the regulatory-based modification disclosure requirement and, if it ultimately opts to retain it, limit the application to AI systems that have been specifically modified or uniquely configured at the request of a non-US government in a manner that materially deviates from its standard commercial

configuration. Standard commercial configurations designed for broad, multi-jurisdictional compliance should not require disclosure under this provision.

Recommendation: **GSA should delete the regulatory-based modification disclosure requirement. If it retains the provision, GSA should limit the application to AI systems that have been specifically modified or uniquely configured at the request of a non-US government in a manner that materially deviates from its standard commercial configuration. It should also issue guidance on how contractors can satisfy this more narrowly tailored obligation.**

VII. Revise Compliance and Reporting Requirements

GSA should also streamline the compliance and reporting requirements, particularly with respect to the disclosure of LLMs used and the incident reporting requirements.

Section (f)(1) requires contractors to disclose all LLMs used or made available in performance of the contract. However, cutting-edge systems, including those commercially available and contracted by the federal government, are highly complex and integrate many components that may be sourced from different vendors and partners. This includes open source models fully controlled and operated by the service provider. Additionally, software and system developers may utilize a range of different tools to conduct research, experiment, and generate and test code that is eventually used in a commercial system – without being directly used in the operation of the system pursuant to a contract. As a result, a comprehensive accounting of all AI systems used in the creation and operation may be disproportionately burdensome, while providing little useful information to federal agencies regarding system performance. Although the proposed revisions refer to FedRAMP are helpful, further alignment is warranted to make the requirement practicable.

In addition, the incident reporting obligations are overbroad and impractical. For example, the threshold for incident reporting includes any incident affecting contractors handling government data, which appears to require reporting even if the contractor has confirmed that the incident has not affected the customer's actual government data. Moreover, daily status updates until resolution are overly burdensome because updates are rarely available in this time frame, and they take away from actual investigation and remediation. As an alternative, we suggest requiring periodic updates until resolution as necessary. Further, Section (d)(2)'s 72-hour notification requirement for known non-adherence to the clause is broader than a typical security incident or breach notification obligation and appears to require reporting any known instance of noncompliance with any requirement in the GSAR clause. Such an overbroad requirement creates a significant administrative burden, whose reporting processes are triggered by identifiable, concrete, and narrowly tailored thresholds that produce risks of specified outcomes. Importantly, GSA should clarify that the 72-hour notification clock begins upon a confirmed determination that a reportable incident has occurred – consistent with the standard applied under FISMA and FedRAMP – and not upon initial anomaly detection or the opening of an investigation.

GSA should also revise the documentation and disclosure requirements. Specifically, Section (f)(7) requires contractors to provide "existing commercial documentation or disclosures" to demonstrate compliance. We recommend that GSA limit this obligation to "as commercially available or reasonably known to the contractor, and tailored to the specific circumstances." In many cases, information may be either unavailable or inefficient to produce. For example, contractors may not have information about specific models, which, in many cases, will be selected for use by the government, and the contractor will not have visibility into the use case.

Recommendations:

- **Revise Section (f)(1) to specify that it does not apply to products and services that have obtained an authorization under FedRAMP/IL5/FISMA, which already disclose components within their authorization boundaries. For systems without such authorizations, the requirement should be narrowed to apply only to the AI models and systems primarily responsible for generating the contract's outputs and should not extend to internal development tools or research processes used prior to contract performance.**
- **For purposes of the disclosure in Section (f)(1), "made available" should mean "accessible to end users through embedded apps, dropdowns, or search bars."**
- **Narrow the incident reporting threshold "affecting any contractor" to prevent overbroad application.**
- **Delete the requirement for daily post-incident status reports, referring instead to periodic updates until resolution.**
- **Narrow the Section (d)(2) non-adherence reporting requirement to specified circumstances, such as security incidents or breaches.**
- **Revise Section (f)(7) to add "as commercially available or reasonably known to the contractor, and tailored to the specific circumstances" to the reference to commercial documentation.**
- **Add an express carve-out for emergency security patches and updates, with post-incident notification.**
- **Clarify that deletion requirements do not override legitimate retention in security logs, audit records, telemetry, backup systems, or FedRAMP-required records.**

VIII. Revise Change Notification Requirements and Data Portability Requirements

Section (i) requires contractors to provide written notice to the Contracting Officer at least thirty (30) calendar days prior to any planned material change affecting services under the contract. As drafted, this requirement is not consistent with standard commercial practice and creates significant operational challenges for contractors offering commercial AI platforms.

Modern cloud and AI services are continuously updated. Requiring thirty days of advance written notice to an individual Contracting Officer for every material change is not scalable across a global commercial customer base and would effectively require contractors to offer a non-commercial, bespoke government-specific product version, contrary to the Administration's direction on commercial terms. More critically, the requirement as drafted would directly conflict with cybersecurity best practices: emergency security patches must often be deployed immediately upon discovery of a vulnerability, and a mandatory thirty-day hold on security updates could expose government systems to known risks during the notice period.

GSA should revise this provision to: (1) define "material change" narrowly to include only changes that adversely affect security posture, relied-upon functionality, or government data handling — excluding routine feature updates, performance improvements, and commercially available configuration changes; (2) exempt emergency security updates and patches from advance notice requirements, substituting a prompt after-the-fact notification obligation to the Contracting Officer; and (3) permit contractors to satisfy the notice obligation through standard commercial release-note, change-management, or trust-center communication channels, rather than individualized written notice to each Contracting Officer.

In addition, Section (h) requires contractors to enable data portability. For SaaS services, users may not be able to export their entire instance in a way that enables them to use it seamlessly in another service. From an operational perspective, that would be extremely difficult, particularly given the broad definition of government data. Indeed, treating a SaaS platform's integrated capabilities as portable could require a provider to transfer, disclose, recreate, or make interoperable proprietary technology, or to enable a functional replica on a competing service. That approach is neither commercially reasonable nor technically

feasible. As a result, GSA should revise the requirement to focus on government data and government-specific, separable deliverables that can be exported in documented, machine-readable formats and clarify that contractors are not required to transfer, disclose, recreate, or make interoperable their pre-existing or independently developed proprietary platforms, frameworks, generalized services, source code, model weights, or other non-separable capabilities. The clause should also acknowledge that the receiving environment may require reasonable adaptation to use exported data or configurations.

Recommendation:

- **Revise Section (i) to define "material change" narrowly (limited to changes adversely affecting security posture, relied-upon functionality, or data handling).**
- **Limit change-notice obligations to controller-controlled changes that materially reduce agreed protections for government data.**
- **Preserve operational flexibility for emergency and unplanned changes.**
- **Revise Section (h) to add a reference to using "commercially reasonable efforts" to provide feasible portability functionality. Limit the portability requirement to government data and separable government-specific deliverables, rather than proprietary SaaS platforms, frameworks, generalized services, source code, model weights, or other non-separable capabilities. Acknowledge that the receiving environment may require reasonable adaptation to use exported data or configurations.**

***

BSA looks forward to continuing to work with GSA and other Administration stakeholders to streamline and accelerate federal acquisition and use of AI to enhance the federal government's ability to perform its important mission and deliver services to American citizens.

Sincerely,

Shaundra Watson
Senior Director, Policy

Jessica Salmoraghi
Senior Director, IT Modernization & Procurement