

March 31, 2026  
General Services Administration  
1800 F St NW  
Washington, DC 20405

## **Re: Response to GSA Draft AI Procurement Clause: "Basic Safeguarding of Artificial Intelligence Systems"**

The undersigned civil rights advocacy and civil society organizations appreciate the opportunity to submit comments in response to the General Services Administration's (GSA) draft AI procurement clause titled "Basic Safeguarding of Artificial Intelligence Systems." We commend GSA for seeking public comment on the proposed clause and we hope our comments will help strengthen and expand the proposed procurement rules.

### **I. Introduction and Scope**

We submit these comments through a civil and human rights lens, grounded in our longstanding work to ensure that AI systems do not replicate, amplify, or launder discrimination against communities already subject to systemic exclusion.

The proposed clause applies broadly to two categories of affected parties:

- **Contractors:** Government contractors providing AI systems or using AI in contract performance.
- **Service Providers:** Entities that directly or indirectly provide, operate, or license AI systems to the government.

Critically, the clause applies to all AI systems used in contract performance, not just those sold directly to the government. This also includes actors who serve as vendors in the supply pipeline as well as AI agents.

This broad scope is significant and warrants careful attention to ensure it does not inadvertently suppress responsible governance practices that exist to protect the public.

### **II. Provisions We Commend**

We recognize and affirm several provisions in the draft that reflect sound governance principles and merit retention in the final rule.

#### **A. Human Oversight, Intervention, and Traceability**

The requirement that contractors *provide a means for the government to implement human oversight, intervention, and traceability* is a strong and necessary provision. As federal agencies continue to integrate AI into high-stakes decision-making processes,

including but not limited to housing, benefits, lending, and law enforcement, the absence of meaningful human control creates unacceptable risk of automated harm. We strongly support this provision and urge GSA to ensure that human oversight mechanisms are designed to be accessible, well-resourced, and exercisable in practice, not merely nominal.

The draft's additional requirement that contractors design AI systems with government monitoring and control capabilities is likewise welcome. Monitoring without the capacity to intervene is insufficient; both capabilities must be mandatory and technically enforced. These and similar provisions are also important to ensure that taxpayer dollars are not wasted on systems that do not work or that cause harm.

## **B. Feedback Mechanisms and Formal Reporting Channels**

The draft's requirement that contractors *establish formal channels for government feedback and performance improvement requests* is equally commendable. This provision operationalizes accountability by creating structured pathways for the government and by extension, the public, to surface concerns about AI system performance without requiring formal incident classification as a prerequisite. We recommend that this provision be retained in the final rule and that GSA consider extending the feedback mechanism to include impacted communities and civil society organizations as formal participants in the feedback loop.

## **III. Recommended Modifications**

While we affirm the positive provisions above, the following changes are necessary to strengthen the draft, address emerging governance gaps, and ensure that the final rule does not undermine civil and human rights protections. Moreover, the changes will help ensure that the public is protected.

### **1. Refine the "Any Lawful Government Purpose" Clause**

#### **The Current Draft**

The draft requires contractors grant the government an irrevocable license to use AI systems for "any lawful government purpose" and further specifies that an AI system "must not refuse to produce data outputs or conduct analyses based on the Contractor's or Service Provider's discretionary policies."

#### **Why This Requires Revision**

The "any lawful government purpose" framing creates a critical governance gap: there are AI use cases, including but not limited to use in surveillance, predictive policing, automated eligibility determinations, and housing allocation, to which existing laws do not yet fully apply or provide inadequate protections.

A blanket "any lawful government purpose" license effectively permits uses that are harmful but not yet explicitly prohibited by statute, for several reasons. The first problem is the lack of legal rules regulating the government's purchases of Americans' sensitive

data—the raw material often fed into AI systems. On top of that, governmental agencies often take an extremely narrow view of constitutional protections, insisting that decisions like *Carpenter v. United States*, 585 U.S. 296 (2018), do not restrict their ability to conduct surveillance that reveals a comprehensive picture of a person’s movements or pattern of life. Finally, existing laws regarding privacy or surveillance are filled with loopholes that the government has long exploited to justify sweeping surveillance and are notoriously difficult to enforce. For all these reasons, specific AI guardrails are warranted, and broad permission to use AI for all “lawful” purposes could open the door to untold harms, including AI-powered mass domestic surveillance.

Moreover, the requirement that AI systems not refuse outputs based on "discretionary policies" conflates responsible safety guardrails with arbitrary refusals. Contractor and Service Provider safety policies including but not limited to those governing harmful content, demographic profiling, and discriminatory inference serve important public interests that the draft should not override. In some instances, as a result of a company’s own assessment and testing, the company itself has designated the capabilities and limitations of a system they create, including indicating appropriate and inappropriate use cases. Federal procurement policies should not ride roughshod over those determinations, especially where a company’s own evaluation shows that a system is not fit for a particular purpose or would result in harm because of biased or inaccurate decisions.

### **Recommended Change**

We recommend two modifications:

1. **Enumerate prohibited use cases explicitly.** The final clause should catalog specific AI use cases that the government will categorically prohibit including but not limited to uses that enable discriminatory targeting, mass surveillance, or automated denial of government benefits without human review. This prevents blank reliance on a residual "any lawful" standard. The final rule should also provide for a process by which additional future use cases that should be prohibited are identified and designated as such.
2. **Preserve rights-protective guardrails.** The final rule should clarify that AI systems may and must refuse to produce outputs or conduct analyses where doing so would violate constitutional or statutory law, including civil rights law, fair housing law, or other anti-discrimination protections and that such refusals are not "discretionary policies" subject to override by this clause.
3. **Permit concrete use limitations.** In addition to the categorical prohibitions included in the final clause, developers should be permitted to provide specific limitations on the use of their models on a case-by-case basis. Because AI developers are often in the best position to assess the limitations of their models, federal contracts should permit developers to identify additional use cases that are not permitted due to technical limitations of their models. We appreciate that federal agencies and contractors alike benefit from specificity, so contractors should seek to identify these limitations as concretely as possible and avoid “discretionary policies” where possible.

## 2. Reconsider the “Woke AI” Provision

### The Current Draft

The draft clause incorporates the substantive provisions of the President Executive Order “Preventing Woke AI in the Federal Government,” 90 Fed. Reg. 35389 (July 28, 2025), and requires AI companies to provide “a neutral, non-partisan tool that does not manipulate responses in favor of ideological dogmas such as Diversity, Equity, and Inclusion,” and prohibits contractors from “intentionally encod[ing] partisan or ideological judgments” into AI outputs.

### Why This Requires Revision

We acknowledge the legitimate goal of preventing AI systems from being weaponized for partisan purposes. However, the current framing is technically flawed and normatively problematic for several reasons.

First, *the question of how to ensure AI system outputs align with specific sets of human preferences is an open question..* Current research demonstrates that reliable behavioral alignment remains an open scientific problem as well as a fundamentally sociotechnical problem requiring research and collaboration between interdisciplinary experts. Mandating “neutrality” without specifying clear standards sets contractors up for compliance failure and may reward performative rather than genuine alignment.

Second, the terms used by the clause, like the Executive Order it follows, are unworkably vague. Terms like “ideological dogmas,” “ideological or partisan judgments,” and “ideological content” are so broad as to be meaningless and unworkable. For example, one dictionary defines “ideological” as “based on or relating to a particular set of ideas or beliefs” and “partisan” as “someone who supports a person, principle, or political party.” Cambridge Dictionary, Ideological Definition; id., Partisan Definition. Read plainly, this would prevent federal AI models from not only reflecting broad policy platforms, such as the President’s own “MAGA,” “MAHA,” or “America First” platforms but even basic tenets of American republican democracy, such as free speech, due process, or equality before the law. Each of those are fairly described as “partisan” or “ideological.”

Likewise, the clause contains no definition of “Diversity, Equity, and Inclusion,” providing no clarity to contractors as to what the clause is meant to prohibit and how to comply. Moreover, “Diversity, Equity and Inclusion” is commonly understood to include many practices that are both lawful and legally required under existing civil rights laws. Broadly characterizing “Diversity, Equity and Inclusion” as “ideological dogmas” misreads the legal and empirical foundations of civil rights law. Disparate impact analysis, for example, is a judicially recognized and congressionally codified tool for identifying discrimination. Federal compliance with longstanding civil rights law — even that prohibiting outright disparate treatment — may require assessing how AI models address protected characteristics and their proxies. That is not an “ideological dogma” but legal compliance. In contrast, an AI system designed to ignore protected class considerations in contexts where law requires their consideration is not neutral — it is potentially violating federal law.

Third, alignment at the behavioral (inference) level as specified in the draft contradicts the need for alignment at the model level. These are distinct technical layers, and the draft's instruction that the requirement "must not be construed to require retraining of the model or alteration of model weights" makes the neutrality mandate technically unenforceable as written.

### Recommended Change

We recommend:

- **Drop the Unbiased Principle as currently framed.** The “woke AI” provision, as written, is technically unworkable and legally unenforceable, and socially regressive.
- **Fund alignment research.** The government should invest in and incentivize mathematically grounded, empirically validated alignment methods at both the model and system levels—approaches contractors can implement and agencies can audit. Alignment remains an open sociotechnical problem; ensuring AI outputs reflect human values requires sustained collaboration across technical, legal, and social domains. Imposing standards such as “neutrality” without building the research infrastructure to define, measure, and audit them risks compliance failure and performative alignment. Federal investment should therefore support not only engineering solutions, but also the normative and empirical research needed to operationalize alignment in high-stakes public contexts.
- **Replace with a civil rights compliance standard.** Contractors should be required to demonstrate that their AI systems, as deployed for government use, comply with all applicable civil rights laws, including but not limited to the Fair Housing Act, the Equal Credit Opportunity Act, Title VI, Title VII, and other statutes like the American with Disabilities Act (ADA).

## 3. Recalibrate the Non-U.S. Regulatory Disclosure Requirement

### The Current Draft

The draft requires contractors to disclose whether AI systems used under the contract have been "modified or configured to comply with any non-U.S. federal government or commercial compliance or regulatory framework." This is consequential for the European Union's Digital Services Act (DSA) and AI Act, as well as U.S. state AI laws.

### Why This Requires Careful Revision

We acknowledge the national security rationale underlying this provision. However, we note several significant risks that the current framing creates.

The current Administration has attempted to significantly roll back foundational civil rights principles that have historically served as constraints on AI systems operating in high-stakes sectors, such as the disparate impact doctrine and the importance of ensuring that systems are built to be inclusive and reflective of representative data.. Several international jurisdictions, including Canada and the European Union, have stepped into this governance vacuum. American AI companies operating in Europe are

required to comply with the EU AI Act and DSA, which impose oversight obligations, transparency requirements, and prohibitions on certain high-risk AI applications.

A disclosure requirement that effectively flags EU AI Act compliance as a potential concern signals to contractors that rights-protective governance adjustments may be treated as a liability in federal procurement. This creates a perverse incentive: contractors may be motivated to strip out civil rights-aligned features to avoid disclosure obligations, the direct opposite of the public interest outcome.

Additionally, mandatory disclosure of regulatory compliance configurations could trigger retaliatory disclosure requirements by trading partners, complicating the competitive position of U.S. AI companies in global markets, an outcome presumably contrary to the Administration's stated goal of winning the AI competition.

Moreover, the disclosure is underinclusive: federal agencies may seek to procure custom models to meet specific needs or direct modifications to models under their existing regulatory frameworks. Other federal agencies in the future may desire to know if current models reflect those changes. Thus, there is an important government purpose not just in disclosure of changes made for non-federal governmental regulatory requirements but also in changes in response to federal requests.

### Recommended Change

We recommend:

- **Expand the disclosure obligation:** Disclosure should be expanded to include changes directed by federal agencies and to include changes made due to compliance requirements and procurement contracts.
- **Create a safe harbor for civil rights-aligned configurations.** The final rule should explicitly exclude from disclosure any AI system configurations implemented to comply with anti-discrimination law, human rights standards, or safety requirements, whether domestic or international.
- **Conduct a proportionality analysis.** GSA should assess whether the disclosure requirement, as drafted, is likely to induce contractors to remove rights-protective guardrails, and revise accordingly.

## 4. Reconsider the Prohibition on Non-U.S. AI Systems

### The Current Draft

The draft bars federal contractors from using any non-U.S.-made AI systems in contract performance, including "any AI components manufactured, developed, or controlled by non-U.S. entities."

### Why This Requires Revision

We appreciate the desire to develop measures to address national security risks associated with foreign AI systems. However, the current prohibition is overbroad and creates significant implementation risks.

The definition of "American AI Systems" as systems "developed and produced in the United States" does not account for the deeply international nature of the AI development ecosystem. Open-source components, training datasets, and algorithmic architectures frequently involve contributions from researchers and engineers in allied nations. A literal application of this prohibition could disqualify large portions of the U.S. AI industry and force reliance on a narrower, potentially less capable, set of vendors.

Furthermore, the prohibition does not address the more pressing risk: the opacity of AI systems, regardless of national origin. A poorly governed, non-transparent AI system developed entirely in the United States may pose greater civil and human rights risks than a well-governed international system subject to rigorous oversight requirements.

### Recommended Change

We recommend:

- **Adopt a risk-tiered approach.** Restrict non-U.S. AI components based on demonstrated security rather than applying a blanket prohibition based solely on geography.
- **Prioritize transparency over geography.** Require all AI systems used in federal procurement regardless of national origin to meet minimum transparency, explainability, and auditability standards.
- **Engage allied nations.** Work with allied governments to develop shared AI governance standards that enable trusted international collaboration without compromising national security.

## 5. Strengthen Data Segregation and Training Prohibition Provisions

### The Current Draft

The draft requires government data to be logically segregated from non-government customers, imposes "eyes off" data handling procedures, and prohibits the use of government data for AI training or fine-tuning.

### Assessment

We support these provisions as written and recommend their retention in the final rule. Data segregation and training prohibitions are essential to prevent the laundering of government data into commercial AI products, which could jeopardize citizens' privacy, who expect their personal information collected by the government to be used for governmental purposes — not enriching private companies. Moreover, private use of federal data without additional safeguards could entrench biases drawn from government datasets including datasets that reflect historical patterns of discrimination in areas such as housing, lending, and law enforcement.

### Recommended Enhancement

In addition to retaining the existing data provisions, we recommend:

- **Extend protections to sensitive demographic data.** The final rule should explicitly prohibit the use of protected classes as inputs to AI system outputs in

contexts governed by civil rights law, unless required for legally authorized purposes such as affirmative marketing or discrimination auditing.

- **Require civil rights compliance auditing as a condition of contract.** Contractors should be required to conduct and disclose third-party audits of civil rights compliance of AI systems used in federal contract performance, with results made available to the contracting agency and, where appropriate, to the public.

## 6. Address the Surveillance Precedent Risk to Americans

### The Risk the Draft Creates

The draft clause combines an irrevocable “any lawful government purpose” license with a prohibition on AI systems refusing outputs based on a contractor or Service Provider’s “discretionary policies.” This combination creates the structural conditions for AI surveillance at scale. By stripping the inference-level guardrails that currently constrain how AI systems may be queried and for what ends, if passed without modification, the clause could authorize the repurposing of federal AI systems for population-level monitoring of Americans without any additional legal authorization.

This risk is compounded by a well-documented historical pattern: federal procurement standards shape private industry norms. When the government normalizes particular AI capabilities and use cases through contract terms, those terms migrate into commercial practice often without the constitutional guardrails that nominally constrain government actors. A clause that licenses AI systems for “any lawful government purpose” without enumerating surveillance prohibitions provides a legitimizing blueprint that private employers, landlords, lenders, and insurers can cite to justify deploying the same systems against the Americans they serve.

### Specific Surveillance Use Cases of Concern

Without explicit prohibitions, the draft clause would leave Americans vulnerable to AI-enabled surveillance across several high-risk domains that we have identified through our civil rights work, that include but are not limited to:

- **Tenant screening and housing applicant profiling.** AI systems procured for federal housing programs could be repurposed by private landlords and screening companies to conduct behavioral profiling of prospective tenants. This includes inferring risk scores from social media activity, purchasing patterns, or geographic mobility in ways that proxy for race, national origin, and familial status in violation of the Fair Housing Act.
- **Predictive policing and neighborhood risk scoring.** Federal AI systems used in law enforcement contexts could be repurposed to generate neighborhood-level risk scores that direct disproportionate surveillance toward communities of color thereby encoding patterns of historical over-policing into automated systems that operate without meaningful judicial oversight or Fourth Amendment constraint.
- **Benefits eligibility surveillance.** AI systems deployed for housing voucher administration, public assistance determinations, or disability benefit reviews

could be configured to conduct continuous monitoring of recipients thereby flagging behavioral anomalies or spending patterns as grounds for automated disqualification without human review or due process, in violation of both the Privacy Act of 1974 and fundamental principles of administrative fairness.

- **Immigration status inference from behavioral data.** Federal AI systems with access to government databases could be repurposed to infer immigration status from behavioral signals including but not limited to language use, remittance patterns, geographic movement, or community association and share those inferences with private employers, landlords, or financial institutions, chilling constitutionally protected activity and exposing vulnerable communities to economic and physical harm.

### Recommended Change

We recommend three targeted modifications to address the surveillance precedent risk:

- **Conduct a private-sector spillover impact assessment before finalizing.** Before the final clause is incorporated into GSA Multiple Award Schedule contracts, GSA should commission an independent assessment of the likely spillover effects of the clause's use-right framework on private-sector AI deployment practices. The assessment should specifically evaluate whether the "any lawful government purpose" license and the prohibition on inference-level refusals are likely to be cited by private actors as justification for deploying equivalent capabilities against Americans in unregulated commercial contexts, including housing, employment, credit, and insurance.
- **Require contractors to certify that AI systems will not be repurposed for surveillance.** Contractors and Service Providers should be required to certify, as a condition of contract, that AI systems provided to or used in support of the government will not be repurposed by themselves, their affiliates, or downstream commercial licensees for tenant screening, benefits recipient surveillance, predictive policing, or immigration status inference without explicit legal authorization and appropriate civil rights safeguards. This certification should be enforceable, with meaningful consequences for breach.
- **Create a public notice requirement for AI monitoring and surveillance.** Whenever a federal agency deploys an AI system procured under a GSA contract for monitoring or surveillance of individuals or communities, the agency should be required to publish advance public notice that: (a) identifies the purpose and scope of the monitoring or surveillance; (b) describes the AI system and data sources involved; (c) identifies the legal authority for the monitoring or surveillance; and (d) specifies the civil rights and privacy safeguards in place. This requirement would operationalize transparency as a structural check on surveillance mission creep and help give affected communities notice and the opportunity to challenge improper use.

## IV. Civil Rights as a Structural Governance Framework

We urge GSA to recognize that civil and human rights frameworks are not merely one policy consideration among many in AI governance. They are the structural foundation without which AI governance cannot be trusted, legitimate, or durable.

The draft clause, as currently written, contains important elements of structural accountability, particularly the human oversight and feedback mechanism provisions, but also contains provisions that, if adopted as drafted, would weaken the civil rights-protective governance features that many AI systems must be required to implement.

GSA has a unique opportunity, through this procurement clause, to signal that the federal government's use of AI will be governed by the same principles of equal protection, due process, and non-discrimination that govern federal action more broadly. We urge GSA to seize that opportunity. At minimum, the references to “known biases” and “bias” throughout the clause should expressly include discrimination that violates federal civil rights laws.

## V. Conclusion and Summary of Recommendations

NFHA respectfully urges GSA to adopt the following modifications to the draft clause before finalizing and incorporating it into GSA Multiple Award Schedule contracts:

1. **Refine the “Any Lawful Government Purpose” Standard:** Enumerate prohibited AI use cases explicitly and preserve rights-protective safety guardrails from override.
2. **Revise the “Woke AI” Provision:** Drop the Unbiased Principle as currently framed; fund model-level alignment research; replace with a civil rights compliance standard.
4. **Recalibrate Non-U.S. Disclosure:** Narrow the disclosure obligation to genuine national security risks; create a safe harbor for civil rights-aligned system configurations.
5. **Refine the Non-U.S. AI Prohibition:** Adopt a risk-tiered approach based on demonstrated security risk rather than a blanket geographic prohibition.
6. **Strengthen Data Protections:** Retain and enhance data segregation and training prohibitions; add demographic data protections and mandatory bias auditing.
7. **Retain Human Oversight and Feedback Provisions:** These are strong provisions and should be retained and strengthened in the final rule.
8. **Address the Surveillance Precedent Risk:** Conduct a private-sector spillover impact assessment; require contractor certification that AI systems will not be repurposed for surveillance; create a public notice requirement for AI monitoring or surveillance. Anchor these protections in the Fourth Amendment, the Privacy Act of 1974, and state-level AI governance frameworks.

A procurement clause that incorporates our recommendations will set expectations for AI developers, deployers, and government institutions that use AI to serve the public. We urge GSA to update their procurement rules so that AI systems are centered on civil rights principles.

Thank you for considering our recommendations.

Sincerely,

1. National Fair Housing Alliance (NFHA)
2. Government Information Watch
3. The Leadership Conference on Civil and Human Rights
4. United Church of Christ Media Justice Ministry
5. American Civil Liberties Union (ACLU)
6. Color of Change