



General Services Administration  
1800 F St NW  
Washington, DC 20405

April 2, 2026

**Re: Response to GSA Draft AI Procurement Clause: "Basic Safeguarding of Artificial Intelligence Systems" (GSAR 552.239-7001)**

To Whom It May Concern:

We write to express grave concerns about the proposed changes to the solicitation provisions and contract clauses that have been proposed by the General Services Administration as part of 552.239-7001, Basic Safeguarding of Artificial Intelligence Systems, in all future solicitations and contracts for Artificial Intelligence capabilities. Among other things, these changes require any vendor for government contracts to “[grant] to the Government an irrevocable, royalty-free, non-exclusive license to use the AI System for the duration of this contract for any lawful Government purpose” as part of the terms and conditions of the contract.

A standard allowing the government to use AI systems for “any lawful use” may sound reasonable in the abstract, but the phrase conceals two distinct categories for misuse that become problematic when applied to concrete scenarios.

**Lawful Uses That Violate Public Norms and Vendor Safety Policies**

As evidenced by the recent dispute between Anthropic and the U.S. Department of War, there are myriad government uses that are clearly legal but conflict with AI industry terms of service and the public’s expectations as to how emerging technologies will be responsibly integrated into state functions. The acceptable use policies, and internal safety standards, of every major frontier AI lab restrict applications in ways that could be violated by any of the following scenarios in which the government is likely to be operating lawfully:

*Psychological profiling of benefits applicants.* A federal agency could use procured AI to analyze the language, tone, and lifestyle patterns of disability or welfare applicants to build psychological profiles to optimize denial rates for those “most likely” to abuse such benefits or reduce administrative overhead by targeting those most likely to abandon an appeal for claims denials. Nothing in current law prohibits profiling applicants for these purposes. Yet the terms of service of most AI companies prohibit using their models to manipulate, exploit, or psychologically profile vulnerable populations. Under an “any lawful use” regime, those protections instantly vanish.

*Surveillance pattern analysis.* Federal law enforcement agencies possess broad authority to monitor publicly available information, social media activity, and retrieve communications metadata. AI could be used lawfully to build behavioral pattern maps of individuals by aggregating and analyzing this publicly available information at a scale and speed that no human analyst could ever hope to match. No wiretaps or court orders would be needed, and no legal remedies would immediately present themselves for those subject to such monitoring. This is precisely the sort of “mass surveillance” that has raised concerns during the Pentagon’s recent conflict with Anthropic. The subsequent public outcry on this subject has been deafening, and such “lawful” usage clearly violates the norms of even those willing to give broad, security-related deference to the state.

*Employee loyalty screening.* Both in matters of recruitment and retention, federal agencies have broad authority to investigate employees for the purposes of adjudicating security clearances and position suitability. AI could be lawfully employed to continuously monitor the communications, social media accounts, and professional networks of federal employees for indicators of “disloyalty” or ideological deviation. One need not even ascribe governmental malice for such activity; security clearance backlogs are legion and those in sensitive national security or homeland security positions should reasonably be expected to be loyal to the United States government. However, such a “lawful use” would create a chilling effect on First Amendment activity that would be technically legal but normatively corrosive.

These hypothetical abuses do not require a rogue actor, operating somewhere deep in the bowels of the federal government, in order to be realized. Each of these scenarios is emblematic of potential good-faith efforts to comply with instructions to optimize state functions through the use of AI with no limitations beyond what is legally permissible. The same vendor guardrails that would be removed by the GSA’s proposed rule - and the public norms on which they are based - exist precisely because legality is a necessary, but ultimately insufficient, standard upon which to base governmental actions related to powerful AI systems.

### **“Lawful” as Defined by the Administration in Power**

Similarly, while the “any lawful use” standard seems stable across and between administrations, it is really only as stable as the legal interpretations of the administration wielding it. Generally speaking, such interpretations tend to expand over time and consistently push outward.

For example, the Obama Administration oversaw the unprecedented expansion of surveillance under Section 702 of the Foreign Intelligence Surveillance Act. As part of this expansion, the administration maintained a classified “disposition matrix” for targeted killings, including of some U.S. persons, and conducted drone strikes where Congress had not yet authorized the military use of force. The legal justifications for both sets of activities was legally contested, but

the programs operated for years before meaningful judicial review occurred<sup>1</sup>. If the AI systems that currently power the Pentagon's intelligence analysis systems had been operating then, one could easily imagine the number of strikes and speed of target identification having increased by orders of magnitude for a set of programs for which both the courts and Congress took issue.

The courts have always served as a bulwark against what is illegal, and that is necessary in its own right. Traditional bureaucracies are slow, and historically that slowness has served as an additional *de facto* safeguard against both injustice and abuse. When an administration pursues a legally questionable policy, the time required for human bureaucrats to implement it at scale naturally provides a window for courts to issue injunctions, congressional hearings to take place, and for public debates to take hold. AI collapses that window entirely. The GSA's proposed changes to contracting solicitations create an opportunity for policies that would have taken months of human processing to implement across the federal government to be executed in hours, days, or weeks. By the time a court is able to rule on the action, the harm is already done. Benefits have been terminated. Investigations have been launched. A policy of enabling "all lawful use" strips away one of the last public safeguards we have against tyranny. We urge you to reconsider adoption of these proposed changes for our federal acquisitions system.

Sincerely,

*Morgan C. Plummer*

Morgan Plummer  
VP of Policy Design & Delivery

*Justin Bullock*

Justin Bullock  
VP of Policy

---

<sup>1</sup> See the FISA Amendments Act of 2008; the DOJ "White Paper" on the lawfulness of targeted killing of a U.S. citizen (2013); reporting on the "disposition matrix" for drone targeting.