OCTANT ADVISORY | POINT OF VIEW

When AI Governance Becomes a Contract Term

What GSA's proposed LLM safeguarding clause signals for the federal AI market, and how to get ready before it lands

By Ian McCulloh and Maria Chaloux, Octant Advisory, June 2026

On June 17, 2026, GSA published a draft acquisition clause, **GSAR 552.239-7001, "Basic Safeguarding of Data within Large Language Model Artificial Intelligence Systems."** The comment window closes August 3, with a public listening session on July 14. On its face it looks like a routine cybersecurity update. Look closer and it's the first government-wide attempt to write AI governance straight into the contract, and the version that becomes final will shape which companies can sell AI to the federal government and on what terms.

We've spent our careers on the leadership and operating-model side of large technology transformations. The security mechanics aren't what should worry contractors here. The structure should. GSA is moving AI governance from a voluntary corporate practice to an enforceable contract obligation, and it's putting most of the resulting risk on the prime contractor, often the party with the least control over how the underlying model behaves.

What the clause does

The clause applies whenever an LLM processes "Government Data," a term defined to cover prompts, source documents, generated responses, analyses, logs, metadata, synthetic data, and derivative information. It can be inserted into the Federal Supply Schedule, GWACs, OASIS+, and commercial-item contracts. It reaches a company even when that company doesn't build a model: integrating a commercial model, a cloud-hosted AI service, a retrieval system, or an agentic workflow is enough to trigger it.

Within that scope it does seven things at once. It gives the Government ownership of its data and bars contractors from using that data to train or improve a model. It imposes "eyes-off" handling, segregation, and certified deletion. It defines four supply-chain roles (developer, system operator, system integrator, service provider) and requires the prime to identify each one and flow requirements down. It mandates disclosure and documentation aligned to the NIST AI Risk Management Framework. It pushes contractors toward U.S.-controlled models and away from adversary-government influence. It sets a 72-hour incident-reporting clock and a 30-day advance-notice rule for material model changes. And it requires LLMs to be "truthful, objective, and politically neutral," subject to Government testing.

Most of it is defensible in principle, but the problem is how the obligations are allocated.

The core problem: accountability without control

Read the supply-chain and change-management provisions together and one issue dominates. A prime that integrates GPT, Claude, or Gemini through Azure or Bedrock is being asked to vouch for how those models behave and who controls them, when the companies behind them won't negotiate bespoke terms with a mid-sized integrator.

The mismatch is structural. The Government sees one solution. The contractor sees a stack: a commercial model, a cloud platform, a retrieval layer, a vector database, a logging tool, a subcontractor application, and a user-facing interface, each with its own vendor and its own contract. The clause makes the prime answer for the whole stack.

The 72-hour incident clock runs through the prime even when the incident starts with a model developer or cloud platform the prime didn't cause and didn't detect first. The 30-day advance-notice requirement, plus a concurrent evaluation window on successor models, assumes the prime can make a frontier lab hold a version or pre-brief a change. Commercial AI providers ship updates on their own cadence. A routine product decision the prime had no say in can put it in breach.

The clause leaves one thing unresolved: the prime's compliance position is only as strong as its least transparent, least cooperative technology provider. The prime carries full accountability with only partial control. That gap is the most likely thing to be reworked in comments, and it's where contractors should focus now.

Where Government data lives

The definition of Government Data is unusually broad, and easy to underestimate. It isn't limited to source documents or personally identifiable information. It reaches prompts, outputs, account information, analyses, logs, metadata, annotations, embeddings, synthetic data, and derivative information. That pulls in material many firms treat as operational exhaust: telemetry, usage history, debugging data, evaluation results, support records, and product analytics.

So the contractor needs a clean answer to a plain question: where does Government data go? The honest answer usually spans production databases, access logs, object storage, vector databases, cache layers, monitoring tools, ticketing systems, backup media, disaster-recovery environments, and subcontractor systems. The clause then asks the contractor to delete all of it at contract end and certify the deletion in writing. That is easy to promise and hard to prove, and the Government will eventually ask not just whether data was deleted but how the contractor knows.

Where the clause is well-aimed

Not all of it is overreach. Three pieces are sound and achievable today. The prohibition on using Government Data to train or improve a model is the right line, and the enterprise and government tiers of the major AI platforms already exclude training contractually. Government ownership of its own data, with certified deletion at contract end, is conventional and overdue for AI systems. And the data-portability requirement, which forces standard formats and bars licensing barriers to export, is the strongest pro-competition language in the document. It reduces vendor lock-in over the life of a contract, which is good for agencies and bad for incumbents who relied on proprietary formats to hold an account. Most contractors will underrate it. That's a mistake.

On model neutrality: a real problem, an unclear compliance path

The most-debated requirement is that models be “truthful, objective, and politically neutral,” tested against Government benchmarks the contractor may not see in advance, with termination-for-cause and decommissioning-cost liability attached. The easy critique writes itself: “neutral” and “ideological content” are hard to define, and two evaluators can disagree about whether an answer is neutral or sufficiently uncertain. That critique is fair, and incomplete. The requirement is pointing at something real.

Today’s foundation models carry alignment choices baked into their weights. Platforms tune those weights for safety, but the line between safety and viewpoint restriction is thin, and the result is a model whose embedded positions are opaque even to the people deploying it. For government work that’s a genuine problem: an agency can’t audit, explain, or answer to elected officials for a viewpoint it can’t see. GSA’s neutrality language is a blunt instrument aimed at a legitimate target.

The useful response to the clause isn’t to argue the standard away. It’s to push for an architecture that can meet it: keep the model substrate as neutral as possible, and move alignment and policy guardrails into an explicit layer above the model, where the applied rules are transparent, auditable, and accountable to the officials who own the mission. Amazon’s Nova Forge is one step in that direction, separating the substrate from the alignment that sits on top of it. That is how a contractor complies with a neutrality requirement without pretending the underlying model is value-free.

The requirement also isn’t going away. It traces directly to **OMB Memorandum M-26-04, “Increasing Public Trust in Artificial Intelligence Through Unbiased AI Principles”** (December 2025), so it’s policy-driven from above. The market would be better served by working out how to comply, and how to sharpen the language, than by treating the provision as a drafting error. Right now there is more criticism than there are proposed solutions. That gap is worth filling before August 3.

Who wins and who loses

The incentives point to clear market effects.

Advantaged: large primes, vertically integrated technology companies, providers with dedicated government-cloud offerings, and firms that own or directly control their models. They can document their stack, absorb the compliance overhead, and negotiate federal-grade terms.

Squeezed: small and mid-sized contractors that depend on commercial AI platforms but lack the bargaining power to obtain special terms, detailed documentation, or advance notice of changes. They often build strong federal solutions by combining their own expertise with commercial tools. That still works, but the paperwork and engineering discipline around it has to improve.

Unless GSA adds standardized attestations, an approved-provider safe harbor, or compliance obligations tiered to the prime’s role in the stack, the predictable result is vendor consolidation and higher barriers to entry, the opposite of the competition the portability provision is trying to create. That tension between the clause’s pro-competition and pro-consolidation effects is, to us, the most important thing for GSA to resolve before this becomes final.

The bigger signal: governance as market access

Beyond the specific clause, what GSA is really doing is making AI governance a condition of selling to the government. Showing that an AI solution performs won't be enough anymore. Contractors will also have to show that its data rights, supply chain, security architecture, model behavior, documentation, and change management meet federal requirements. Proof of control becomes part of the product.

This is the shift we see coming across the federal AI market no matter how this particular draft is finalized. Agencies aren't going to relax their expectations on data protection, supply-chain transparency, or model accountability. Companies that treat governance as paperwork to assemble after award will lose to companies that build it into how they pick models and price the work.

What to do before the rule lands

The wrong move is to wait for a final rule. Some version of this is coming, and the firms that can already document their stack and secure federal-grade commitments from their providers will win the early work. None of the moves below needs a final rule to justify the investment, and each takes longer to build than the comment window allows.

Move	What it takes
Inventory LLM use	List every place an LLM touches federal work, including AI features embedded inside commercial tools you may not think of as "AI products."
Map the data path	Trace where Government Data goes: prompts, outputs, logs, metadata, embeddings, retrieval indexes, support records, and backups.
Name every provider	Identify each entity in the chain: model developer, operator, integrator, cloud provider, application provider, and subcontractor.
Pressure-test vendor terms	Check your agreements with OpenAI, Anthropic, Microsoft, Google, AWS, and others against the clause: training exclusion, incident notice, change notice, deletion, audit support, FedRAMP evidence, foreign-ownership disclosure.
Build a deletion procedure	Define how you delete and certify across production systems, logs, vector stores, backups, disaster-recovery, and subcontractor platforms.
Stand up incident escalation	Get a path that pulls facts from downstream vendors fast enough to support a 72-hour Government report.
Put a senior owner on it	Model selection, hosting, and provider terms now carry contractual and financial consequences. Treat this as an enterprise-risk decision, not an engineering one.

Where to focus your comments

This is a request for comments, not a final rule, and GSA has already shown it listens. The current draft is the second version; the January 12, 2026 draft was broader, and GSA narrowed it so the

clause now applies only when an LLM processes Government Data. If you engage before August 3, the goal isn't to weaken the Government's protection of its own data. It's to make the final clause workable enough that a responsible contractor can comply. The provisions worth comment:

- Clarify what counts as Government Data, especially logs, telemetry, synthetic data, embeddings, and derivative information.
- Create standard vendor attestations or a safe harbor for major AI and cloud providers.
- Spell out how a prime complies when a commercial provider won't accept full flowdown terms.
- Define "material model change" so it doesn't capture every routine software update.
- Set realistic deletion requirements for backups, logs, and disaster-recovery systems.
- Spell out how neutrality is tested, and favor architectures that keep alignment in an explicit, auditable layer above the model rather than baked into the weights.
- Clarify how Government ownership of custom developments interacts with contractor IP, reusable workflows, prompts, and configurations.

Bring specifics. The notice terms your providers offer, the real mechanics of deletion across vector stores and backups, the concrete cost of a dedicated environment. Comments grounded in operational detail move rules. Vague objections get filed and forgotten.

Sources

GSA, "General Services Acquisition Regulation; Acquisition of Information and Communication Technology; Notice of Listening Sessions and Request for Comments," Notice-MVAC-2026-01, Federal Register, June 17, 2026. federalregister.gov

OMB Memorandum M-26-04, "Increasing Public Trust in Artificial Intelligence Through Unbiased AI Principles," December 2025. whitehouse.gov

About Octant Advisory

Octant Advisory helps organizations convert AI ambition into measurable performance. We work with leadership teams on the governance, workforce, and operating-model changes that make AI investment pay off, the issues that sit underneath clauses like this one. Ian McCulloh built and led a 1,200-person, $1.5B federal AI practice and directs AI executive education at Johns Hopkins. Maria Chaloux built the leadership team behind Accenture Federal Services' growth from $1B to $5.5B over a decade, and spent two decades helping organizations identify and develop the leaders who drive transformation.

Learn more at octantadvisory.com

This document is for general information and does not constitute legal advice. Contractors should consult counsel on the application of GSAR 552.239-7001 to specific contracts.