GW | LAW
GOVERNMENT PROCUREMENT

UNITED STATES COURT OF FEDERAL CLAIMS · JUDICIAL CONFERENCE

Filing Blind

Ethics and Judgment in the Age of Artificial Intelligence

Jessica Tillipman
Associate Dean for Government Procurement Law · The George Washington University Law School
September 2026

GW | LAW
GOVERNMENT PROCUREMENT

Hidden instructions in a court filing

THE FILING

Self-represented litigant, Connecticut state court.

Tiny white-on-white text inserted into filings.

THE INTENDED READER

Instructions directed at any AI system that might process the document, not at the human reader.

The court was not using AI to evaluate filings; the attempted manipulation did not affect the merits.

THE SANCTION

The court revoked e-filing privileges.

Similar incidents reported in other courts.

GW | LAW
GOVERNMENT PROCUREMENT

Benefits and risks

Benefits already visible

Solicitation and clause review, compliance and oversight automation, risk assessment and prequalification, fraud and collusion detection, reduced government touchpoints.

Corruption risks

Exploitation of AI tools for illicit gain, competitive advantage, or manipulation of the procurement process.

Integrity risks

Technical limitations, opacity, market concentration, and transparency failures that undermine fair outcomes regardless of intent.

New ethical questions

Fabrication, hidden instructions, unauthorized tool use, conflicts embedded in tools, and reliance on machine output in decisions that require human judgment.

Background

What the Government is buying → *the stack.*

How the model works → *the chef.*

Why AI procurement is different → *the risk vectors.*

The legal framework that may already apply → *no AI-specific statute.*

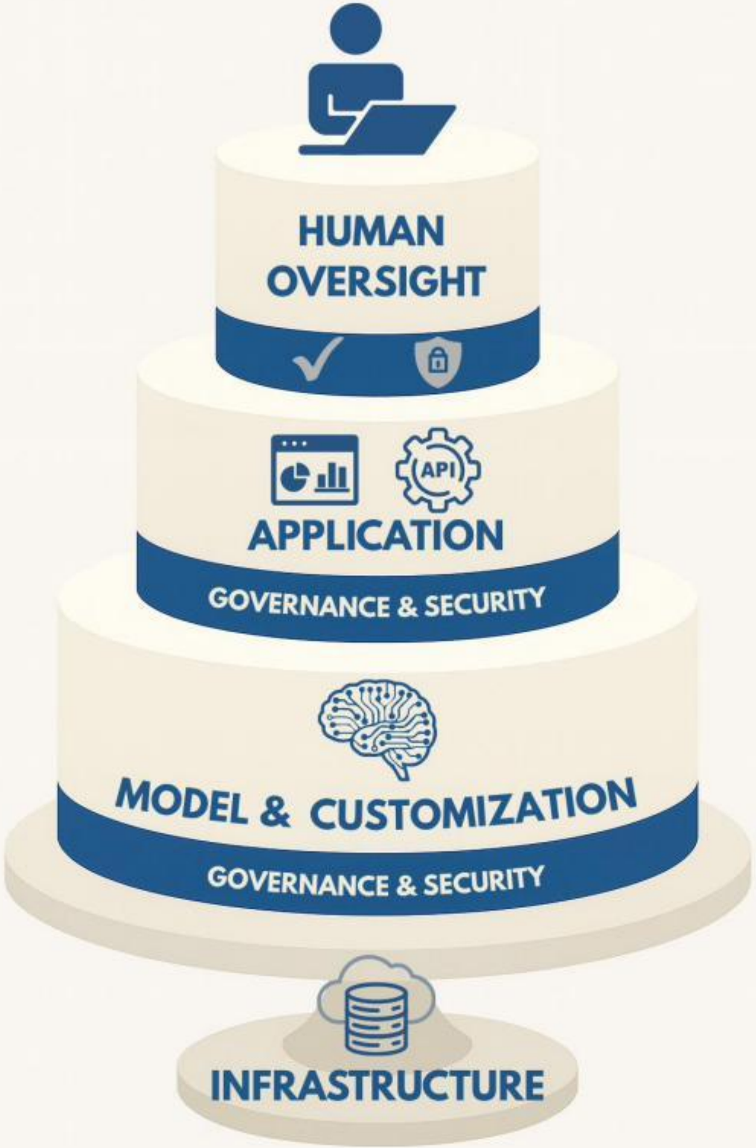

You buy a stack, not an app

A layered supply chain

Every AI purchase sits on infrastructure, a foundation model, customization, and an application — with human oversight on top. Most of it is invisible from the interface the user touches.

Governance wraps every layer

Security, contract terms, and oversight must wrap the whole stack, not only the tier the agency can see.

Concentration at the bottom

Competition at the top; concentration at the bottom. Responsibility extends beyond the button the user pushes.

HOW THE MODEL WORKS

The model is a chef. Her career has five stages.

GW LAW
GOVERNMENT PROCUREMENT

What “we don’t train on your data” leaves out

The word “train”

Undefined unless the contract defines it. The apprenticeship is training too: fine-tuning changes the model itself.

The recipe book

May be silent on where supplied materials are stored, who can access them, and whether they are returned at the end of the contract.

The service records

Silent on the data exhaust (data dust) the vendor retains while the model is in use. The vendor learns from those records without training on them.

Deletion

Deleting the files removes the files. It does not remove what a fine-tuned model learned from them.

Technical characteristics become risk vectors

Market concentration

A small number of leading foundation-model providers and even fewer suppliers of advanced AI chips beneath a varied application layer.

Opacity

Training, architecture, and customization may be difficult or impossible for the agency to inspect fully or independently verify.

Emergent behavior

Systems can exhibit unexpected capabilities and behaviors not captured in their specifications or testing.

Embedded bias

Training, fine-tuning, and design choices can embed biases or preferences the agency did not choose and may not be able to identify.

Hallucination

Confident generation of false content, including citations, capabilities, and past performance.

Appearance of objectivity

Machine output discourages the scrutiny that a human recommendation would receive.

GW LAW
GOVERNMENT PROCUREMENT

The rules that *may* already apply

False Claims Act
A "knowing" false AI-based certification may be a false claim.

Mandatory disclosure
FAR 52.203-13 may apply to AI-enabled misconduct.

Procurement Integrity Act
Potential disclosure of contractor bid or proposal information or source selection information through third-party AI tools.

Organizational conflicts of interest
FAR 9.5 may apply to influence exercised through a tool.

False statements
18 U.S.C. § 1001.

Sherman Act § 1
The risk of shared pricing algorithms.

Truthful cost or pricing data
AI-generated cost data must be accurate, complete, and current.

Suspension and debarment
Present responsibility, including AI governance.

Administrative review
Reasoned decision-making under 28 U.S.C. § 1491(b)(4) and 5 U.S.C. § 706.

No comprehensive AI-specific statute governs these issues.

Integrity issues before the court

Fabrication → *citations, evidence, and records.*

Hidden instructions → *content written for a machine reader.*

Unauthorized tools → *shadow AI and the Procurement Integrity Act.*

Conflicts in tools → *the vendor inside the evaluation.*

Judgment in the record → *the decision the court reviews.*

Fraud risks amplified by AI

The narrative

Fabricated past performance narratives that are realistic, specific, and difficult to detect. Invented projects, invented outcomes, invented references.

The paper

Synthetic documents and certifications: financial statements, compliance attestations, test reports, and personnel qualifications.

The performance

Deepfake oral presentations, reference calls, and facility brochures.

GW | LAW
GOVERNMENT PROCUREMENT

Hallucinated authority is a growing risk

Already happening

Courts are repeatedly encountering fabricated cases, false quotations, and mischaracterized authority in AI-assisted filings.

Why it is different

Generative AI can produce false authority that looks specific, plausible, and professionally written. The error may not be obvious from the face of the document.

Why the risk grows

As AI-assisted drafting becomes routine, more legal work will be generated at greater speed and volume. Without independent verification, that means more opportunities for fabricated material to enter filings and the record.

GW LAW
GOVERNMENT PROCUREMENT

Fabrication suggested in the record

VENERGY GROUP (CBCA)

Audited financial statements produced in discovery were altered: white boxes concealed financial information about overhead costs.

AUGUST 18, 2026

The CEO suggested a former employee with access to his Adobe account and an interest in Grok as a possible source of the alterations. The Board found that Venergy conducted no internal investigation.

The Board found the claim was overstated by more than $1.1 million and dismissed it as a sanction. The termination appeal and the Government claim proceed.

THE RECORD

The risk extends beyond briefs. Generative AI can also fabricate or alter the claims, records, and evidence agencies and tribunals rely on.

GW LAW GOVERNMENT PROCUREMENT

Verification and authentication

1. **Verify** — Confirm material claims outside the submission: past performance databases, personnel records, test reports.
2. **Authenticate** — Confirm key personnel, facilities, and references directly rather than through the offeror.
3. **Require provenance** — Demand source and chain of custody for critical certifications, attestations, and media.
4. **Document** — Record the verification step in the file so the record shows what a human confirmed.

Detection tools produce false positives and false negatives. A flagged result should trigger review, not automatic exclusion.

GW | LAW
GOVERNMENT PROCUREMENT

Content directed at a machine reader

Proposal tailoring / optimization

Drafting toward disclosed evaluation criteria. Repeating the solicitation's terminology because the solicitation values it. Visible to the human reader.

Manipulation/Prompt Injection

Instructions embedded in proposal content that are directed at the AI system rather than the evaluator and designed to alter how the system evaluates the submission.

The distinction turns on whom the content addresses: the evaluator or the evaluation tool.

GW | LAW GOVERNMENT PROCUREMENT

Army prototype RPP, March 2026

The vehicle

Prototype competition under 10 U.S.C. § 4022, issued March 6, 2026.

AI-assisted review

The Government announces that it will use AI to assist in proposal analysis and review, including potentially summarizing content, identifying compliance, and highlighting potential strengths, weaknesses, opportunities, and risks.

Human judgment reserved

Final evaluative judgments and award decisions remain with Government personnel.

Prohibited formatting

Text below 8 point; white-on-white text; text indistinguishable from the background; text hidden in graphics; formatting designed to manipulate or obscure readability. The Government may verify compliance technically and reject noncompliant proposals.

GW LAW
GOVERNMENT PROCUREMENT

When the agent selects the means

INSTRUCTION TO THE AGENT	SIGNIFICANCE
"Insert hidden instructions."	Direct evidence of intent to conceal or manipulate.
"Win by whatever means necessary."	Potential circumstantial evidence of intent.
"Prepare the strongest proposal."	Scienter difficult to establish if the agent selects manipulation on its own.
"Prepare an accurate, compliant proposal; flag unsupported claims for human review."	Compliance evidence, not a safe harbor.

GW | LAW
GOVERNMENT PROCUREMENT

The Procurement Integrity Act and AI tools

The statute

41 U.S.C. § 2102 prohibits knowing disclosure of contractor bid or proposal information or source selection information before award, and any person from knowingly obtaining it.

The tool

Entering protected information into an unauthorized third-party AI service may constitute a disclosure to the provider. Depending on the service and configuration, prompts and outputs may also be retained.

The open question

Whether transmission to an unauthorized AI service constitutes a “disclosure,” and when the provider can be said to “knowingly obtain” the information.

Conflicts of interest at the application layer

GW LAW
GOVERNMENT PROCUREMENT

For issue-spotting, treat the AI application as “Bob”—then ask what you would ask if the vendor were performing the same function directly.

1 **Impaired Objectivity—** Bob evaluates proposals. Bob’s vendor also competes for awards that the application helps evaluate.

2 **Biased Ground Rules —** Bob drafts specifications or evaluation criteria. Bob’s vendor later competes under the requirements the application helped create.

3 **Unequal access —** Bob receives nonpublic, competitively useful agency information. Bob’s vendor may gain access to information that can inform its future competitive strategy.

GW LAW GOVERNMENT PROCUREMENT

Customized-model conflicts: Bob is upstream

The agency contracts with an application vendor, but the application relies on a model customized by someone else.

The customizer

Contractor A customizes a model using its own successful proposals, terminology, and technical approaches as examples of "industry best practices."

The application

Contractor B incorporates that customized model into an agency-facing evaluation tool. The agency sees B's application, not A's upstream design choices.

The competition

Contractor A later competes in a procurement the tool helps evaluate. The model treats A's methods as the benchmark, so its influence appears to evaluators as neutral machine output.

The "Bob" question

If Contractor A were sitting at the evaluation table, the conflict would be obvious. Does FAR 9.5 reach the same influence when A is embedded upstream in the model and reaches the agency through B's application?

GW | LAW
GOVERNMENT PROCUREMENT

Reasoned decision-making: the existing standard

The standard of review

28 U.S.C. § 1491(b)(4), incorporating 5 U.S.C. § 706. State Farm: a rational connection between the facts found and the choice made.

The source selection decision

FAR 15.308: the decision shall represent the source selection authority's independent judgment and shall be documented.

Reliance on others' analysis

The SSA may rely on reports and analyses prepared by others; the record must still support the exercise of independent judgment.

The premise

The framework presumes a human decision-maker who can explain the reasoning in the record.

GW LAW
GOVERNMENT PROCUREMENT

Abdicated decision-making

Rubber-stamped narrative

The SSA adopts the AI-generated evaluation report without independent analysis.

Unstated criteria

Evaluation factors introduced through model drift or hallucination that appear nowhere in the solicitation.

The documentation paradox

More paper and less reasoning. A fifty-page report that no one in the agency can defend.

Questions for the record

Can the source selection authority explain the rationale, not only the rating? Does each finding map to a stated evaluation factor? Does each finding cite the proposal? Did a human make at least one change, disagreement, or tradeoff call?

GW | LAW

GOVERNMENT PROCUREMENT

Questions

Jessica Tillipman

Associate Dean for Government Procurement Law
The George Washington University Law School

jessicatillipman.com