



TECHNICAL BRIEF · BASED ON GAO-26-107681 · GOVERNMENT-WIDE, OMB

OMB AI GUIDANCE FULLY COVERS 2 OF 10 PRIVACY CHALLENGES

2 of 10

EXPERT-IDENTIFIED AI PRIVACY CHALLENGES THAT OMB'S GUIDANCE FULLY ADDRESSED

In March 2026, GAO compared OMB's government-wide AI guidance to 10 expert-identified privacy challenges and found it fully addressed 2. GAO's two recommendations to OMB are open, so primes carry the privacy evidence burden on AI task orders.

THE FINDING

On March 26, 2026, GAO published [Artificial Intelligence: OMB Action Needed to Address Privacy-Related Gaps in Federal Guidance](#) (GAO-26-107681). GAO convened a three-day virtual panel of 12 experts in January 2025, and the panel named 13 key challenges to protecting privacy when agencies use AI. Three were set aside, since GAO found it unreasonable to expect OMB guidance to address them. Against the other 10, GAO compared OMB's government-wide AI guidance, including M-25-21 and M-25-22.

The guidance fully addressed two: building federal workforce skills, and scaling AI systems with privacy protections. It partially addressed the other eight. None went unaddressed. GAO also found the guidance does not specify the types of known privacy-related risks agencies should consider.

GAO made two recommendations to OMB. The first asks OMB to specify examples of known privacy-related risks agencies should consider when updating their AI policies. The second asks OMB to facilitate information sharing or issue guidance on topics that include auditing AI models that contain sensitive information, storing data so sensitive data can be separated, privacy performance metrics, AI-specific privacy impact assessments, and privacy-performance tradeoffs.

WHY IT KEEPS HAPPENING

Process gets covered and mechanics get left out, a pattern GAO documents eight times. M-25-21 requires pre-deployment testing and AI impact assessments for high-impact uses, but GAO found it does not say how to keep sensitive data from being exposed while a model is evaluated. M-25-21 pushes data traceability, but it does not say how to store data so sensitive records can be separated. M-25-22 requires privacy terms in AI contracts and

early involvement of agency privacy officials, but GAO's second recommendation asks OMB for clear privacy rules for agencies developing AI solutions internally.

Experts on GAO's panel explained why the mechanics are hard. Removing data to protect privacy can cost model performance, and a model is difficult to evaluate without sensitive information. Agencies also run privacy impact assessments inconsistently, partly because it is unclear what public input an AI assessment needs.

OMB did not comment on the draft report or supply documentation.

WHY IT MATTERS TO YOUR BID

Until OMB answers, each agency fills the gaps itself, and a solicitation may carry the agency's reading of them. M-25-22 already tells agencies to disclose in solicitations whether a planned AI use meets the high-impact threshold, and it requires contracts to ensure compliance with the minimum practices for those uses. M-25-21's definition of high-impact AI includes outputs that affect privacy.

In a technical volume, that shows up as a request for "understanding of the requirement" on privacy for a named use case. An evaluator can ask how the offeror will test a model without exposing sensitive records, who reviews the privacy impact assessment, and what the offeror will measure if the agency removes fields from a dataset.

TECHNICAL APPROACH

The method below comes from Avalon's [Cloud AI/ML Services](#) offering. It describes how the work is done, at the methodology level, inside the agency's own authorized cloud tenant.

- 1. Score the use case and its risk tier.** Write the use case, its users, and success metrics such as accuracy, precision and recall, or groundedness for generative AI. Score it for data availability, mission value, buildability inside the current cloud boundary, and M-25-21 risk tier. M-25-21 section 5 counts effects on privacy toward high-impact status. The Chief AI Officer (CAIO) makes the determination.
- 2. Assess data readiness source by source.** Inventory each source system, test its quality, and record the access and rights gaps. Document the data-sharing and privacy approvals each source requires, and whether they exist. The data steward supplies this, and the output is a Data Readiness Assessment, not a data platform.
- 3. Map governance artifacts to the minimum practices.** Write the model card, the AI impact assessment inputs, and the monitoring plan, organized by AI RMF function (Govern, Map, Measure, Manage) and the AI 600-1 generative AI profile. M-25-21 section 4(b) says the impact assessment should reference privacy impact assessments, so the inputs point to the privacy impact assessment for the system.
- 4. Evaluate on real agency data.** Run the evaluation against the metrics fixed at kickoff, on representative data instead of curated samples, inside the tenant. The Evaluation Report records test method, results, known limitations, and failure modes. A documented miss is a valid exit.
- 5. Produce security evidence the ISSO can read.** Confine the work to FedRAMP-covered AI services in the agency boundary, and brief the ISSO in week one. Collect control-relevant evidence for the AC, AU, CM, SI, and SA families, including SA-11 developmental testing and SI-4 monitoring. Keep the architecture model-portable, with code, prompts, and documentation owned by the agency, consistent with M-25-22.

Avalon documents. The CAIO decides risk tier and waivers, assessors assess, and the Authorizing Official decides whether anything operates. Avalon does not certify a system as compliant.

WHAT "FIXED" LOOKS LIKE

GAO has not published closure evidence, because both recommendations are open. These are the targets that would let GAO close them, written as targets.

- OMB issues examples of known AI privacy risks that agencies can use when updating policies.
- OMB issues guidance, or runs information sharing through a forum such as the Chief AI Officer Council or Federal Privacy Council, on the topics in recommendation 2.
- Agencies carry AI-specific items in their privacy impact assessments, with a named owner.

At the use-case level, fixed means each minimum practice has documented evidence and a named reviewer, and the impact assessment cites a privacy impact assessment.

STATUS CHECK

OPEN Both recommendations to OMB were listed open on GAO's product page when read on October 8, 2026. OMB did not comment on the report, and GAO said it intends to follow up.

AGENCY	REPORT	RECOMMENDATION	STATUS	AS OF
OMB	GAO-26-107681	1: specify examples of known privacy-related risks	Open	GAO's last update, March 2026
OMB	GAO-26-107681	2: facilitate information sharing or issue guidance on listed privacy topics	Open	GAO's last update, March 2026

TEAMING FIT

NAICS	541512, 541511
AVALON'S SUB ROLE	Use-case risk-tier and feasibility scoring, data readiness assessment, M-25-21 and AI RMF governance artifact mapping, and evaluation on agency data under the prime's AI task order
SMALL-BUSINESS STATUS	Small business; HUBZone (in progress)
LIKELY VEHICLES	GSA MAS (SIN 54151S IT Professional Services); GSA OASIS+

If an agency AI pilot or modernization task order is in your capture plan, Avalon teams as the sub that turns the privacy open questions into a data readiness assessment, governance artifacts mapped to M-25-21, and an evaluation report the agency's CAIO, privacy office, and ISSO can read.

BRANDON IRELAND, FOUNDER · CONTACT@AVALONTECHSERVICES.COM · AVALONTECHSERVICES.COM

This brief analyzes publicly released oversight reports. It does not describe work Avalon performed for any agency named.