

Government AI in Practice

A Research Newsletter from the Government IT/AI Governance Initiative (GIAG)

Issue 2 | Late March 2026 | ThinkCapital LLC | thinkcapital.org

EARLY SIGNAL

An observation surfacing in early GIAG conversations: agencies racing to meet the M-25-21 high-impact AI documentation deadline are treating the classification exercise itself as a governance action. Designating a system as high-impact or non-high-impact is being conflated with managing the risks that designation was meant to surface. The deadline creates a documentation event. Whether it creates a governance change is a different question, and almost no one is measuring the difference.

From the Editor

In our previous edition of this newsletter we examined what the research literature cannot currently answer about NIST AI RMF implementation: how it varies across agency types, where documented compliance diverges from operational accountability, and what meaningful oversight actually looks like when it is working. That issue framed the research questions. This one addresses the policy environment those questions are operating inside.

Issued in April 2025, OMB Memoranda M-25-21 and M-25-22 represent the most concrete federal AI governance policy deliverables in the current cycle. They establish deadlines that are now arriving. The 365-day compliance window for high-impact AI documentation under M-25-21 closes this month. Agencies are reporting compliance. The question worth asking is what that compliance actually consists of.

This issue also examines a question that agency CIOs are raising privately but few are putting in writing: whether the current federal AI governance mandates will produce operational accountability or a well-documented version of the status quo. The answer depends on decisions being made right now, under deadline pressure, with imperfect information.

That is the argument this issue makes. Policy moves at the speed of issuance. Organizations change at the speed of learning. The gap between those two rates is the governance problem nobody has put a measurement framework on. GIAG intends to.

Michael Bragen, Principal, ThinkCapital LLC

ANALYSIS

When Policy Moves Faster Than Organizations Can Learn

The gap between AI policy issuance and implementation fidelity is not a compliance problem. It is a measurement problem. And it is getting larger.

The Policy Velocity Context

Between March 2024 and April 2025, federal AI governance policy moved faster than it had in the preceding decade. OMB Memorandum M-24-10 established agency Chief AI Officer requirements and risk management obligations. M-25-21 and M-25-22, issued April 3, 2025 under Executive Order

14179, rescinded and replaced that framework within twelve months. The result is two binding memoranda governing how federal agencies develop, deploy, and procure AI, with implementation clocks already running.

OMB M-25-21: Accelerating Federal Use of AI through Innovation, Governance, and Public Trust (April 3, 2025)

Directs agencies to appoint Chief AI Officers, establish AI Governance Boards, and develop agency-wide AI strategies. Establishes minimum risk management practices for high-impact AI use cases, including pre-deployment testing, AI Impact Assessments, and ongoing monitoring, with a 365-day documentation deadline. For government IT leadership: the Chief AI Officer now carries binding responsibility for high-impact AI governance. This is a named accountability role with a compliance clock attached to it.

OMB M-25-22: Driving Efficient Acquisition of Artificial Intelligence in Government (April 3, 2025)

Applies to contracts awarded or renewed on or after October 1, 2025. Requires performance-based acquisition, pre-award testing of proposed AI solutions, and contractual provisions covering data ownership, privacy, vendor lock-in protections, and disclosure of unanticipated AI use. GSA was required to publish a procurement tools repository within 200 days of issuance. For government IT leadership: procurement officers are now accountable for ongoing AI contract performance monitoring across the full contract lifecycle, not just at award.

Taken together, the two memos shift the federal posture toward faster deployment with named accountability guardrails. For government IT leadership, that means two things are now happening simultaneously: organizational pressure to move on AI adoption, and new accountability structures that assume the governance capability to manage what gets deployed. Whether that capability exists, and how it varies across agency types, is the question this newsletter is tracking.

***Compliance documentation and operational governance are measuring different things.
Almost no one is tracking the distance between them.***

What M-25-21 Actually Requires

M-25-21 is more operationally specific than its predecessor in several respects. It requires agencies to identify high-impact AI use cases using a defined categorical framework. For those use cases, it mandates minimum risk management practices including pre-deployment testing, AI Impact Assessments, and ongoing monitoring. Agencies must document implementation within 365 days. Chief AI Officers carry named accountability.

The memo also removes a number of reporting requirements from its predecessor, which the policy record characterized as barriers to adoption. The direction is clear: reduce friction, move faster, trust agency leadership to manage risk. The governance problem that creates is that removing reporting requirements also removes visibility. The prior framework generated data. M-25-21 trades that data for speed. Whether speed without visibility produces better governance outcomes is an empirical question. No one has answered it yet.

The Fidelity Gap: What the Deadline Measures

The 365-day compliance window measures one thing: whether agencies have documented that high-impact AI use cases are identified and minimum risk management practices are in place. It does not measure whether those practices are functioning, whether the AI Impact Assessment process changed any deployment decision, or whether ongoing monitoring is occurring at the system level rather than existing as a policy statement.

Without a measurement framework that distinguishes documentation activity from governance effect, agencies and OMB alike are operating without feedback. Policy that cannot be evaluated cannot be improved. Capers Jones spent thirty years documenting what happens when organizations measure activity instead of outcomes: the gap between what management believes and what practitioners know grows until a failure is visible enough that someone asks the right question. In federal AI governance, that question is not whether agencies filed their documentation. It is whether the documentation reflects a changed organizational capability.

Speed without visibility is not a governance posture. It is a bet on capability that has not yet been measured.

The State Agency Dimension

Federal velocity is one variable. State legislative activity is another. More than 1,000 AI-related bills were introduced in state legislatures in 2025, creating a compliance environment for agencies operating across jurisdictions that is genuinely novel. State CIOs describe the pilot-to-production transition as their defining challenge for 2026, per NASCIO. Governance frameworks built for pilot environments were designed under different assumptions about scale, risk surface, and accountability architecture. Scaling a pilot does not scale the governance structure that covered it.

GIAG Stream One is examining this transition directly: do agencies that developed measurement-grounded governance for their pilots produce better operational accountability at production scale than agencies that treated pilots as technology evaluations?

Where This Is Heading

The most likely near-term trajectory for federal AI governance is a compliance surface that looks more robust than the operational reality underneath it. Agencies will file their documentation. AI Governance Boards will be stood up on paper. Chief AI Officers will report. The M-25-21 deadline will be met in the same way that FISMA deadlines were met for two decades: through diligent documentation of practices that exist in policy but are unevenly embedded in operations.

At the state level, the picture is more fragmented. State agencies lack the centralized compliance infrastructure built through FISMA and FedRAMP. They are absorbing AI deployment pressure from above, legislative mandates from their own statehouses, and vendor urgency from below, without a shared measurement vocabulary. The bills introduced in 2025 created requirements but not frameworks for evaluating whether those requirements produce governance outcomes.

The likely result is a two-tier accountability environment. Agencies with measurement discipline and strong Chief AI Officer leadership will develop genuine governance capability. Agencies without that foundation will produce compliant documentation and continue operating largely as before. That performance gap will not appear in any compliance report. Whether it produces observable failures quickly enough to prompt corrective action, and what form those corrections take, remain to be seen.

THE COMPLIANCE GAP

OMB memoranda carry real authority, but M-25-21 and M-25-22 have no defined enforcement mechanism around compliance. There is no OIG audit program structured around the requirements, no published consequence schedule, and no public compliance dashboard. The accountability is named and largely unverified.

This pattern is not new. FISMA was formally binding for years before audits with teeth arrived. Policy authority consistently precedes verification infrastructure. Agency CIOs who have seen this cycle are not wrong to ask whether the April deadline is a governance milestone or a documentation exercise with consequences yet to be specified. That question carries a risk on both sides, and the risks are not symmetric.

The CIO who builds genuine governance infrastructure becomes accountable for what the record shows. An agency that reports accurately about an incomplete system may be more exposed than one that reports minimally about a system it has not examined. In a self-reported compliance environment, candor is not obviously rewarded. The CIO who waits faces a different calculation: if an AI system failure occurs where M-25-21 applies, the absence of documented minimum risk management practices is not a neutral fact. AI failures in government deployments are not a hypothetical. At current adoption scale, they are a timing question. The record agencies are building now is the record that will exist when one occurs.

State CIOs watching the federal compliance pattern through NASCIO are making a rational calculation. Waiting for federal signals before committing is understandable. It is also a governance failure in motion. AI systems are going into production in state agencies today, under frameworks designed for pilots, without oversight infrastructure adequate for the risk surface those systems are operating in.

The most dangerous position is not non-compliance. It is checkbox compliance. An AI Impact Assessment completed because the memo requires one, signed off by someone without authority to change a deployment decision, filed where no operational process references it, is documentation. Its real cost is that it gives leadership a basis for believing the problem has been addressed, removing the urgency to build what is actually missing. Compliance and governance are not the same thing. That distinction is the core measurement problem GIAG is designed to examine. The window for closing the gap before the first visible failure in a high-impact use case is shorter than the compliance calendar suggests.

RESEARCH UPDATE

GIAG is in active participant recruitment. Structured interviews are being scheduled with government IT practitioners across federal, state, and local contexts. Both streams are open for participation. If you work in government AI deployment or oversight and are willing to speak confidentially, details are at thinkcapital.org/research.html.

STREAM ONE

NIST AI RMF Implementation Fidelity in Government Practice. Examining how federal and state agencies are operationalizing the NIST AI Risk Management Framework, distinguishing compliance documentation from genuine operational embedding. With M-25-21's 365-day deadline closing, implementation variation across agency types is becoming observable. Practitioners involved in the high-impact AI designation process are a particular priority.

STREAM TWO

Human Oversight Quality in Agentic AI Deployments. Examining oversight architecture where AI systems act sequentially, chain decisions, or operate within automated workflows. Core question: when a human reviewer is nominally present but structurally positioned after consequential decisions have been made, does that oversight satisfy the functional intent of M-25-21's minimum risk management requirements? Early conversations suggest the answer varies significantly by agency type and context.

Participate in the GIAG Research

Interviews are confidential, conducted remotely, and require approximately 30 minutes. Attribution is optional. Participants receive a summary of preliminary findings before public release.
thinkcapital.org/research.html

WHAT TO WATCH

DEADLINE

M-25-21 High-Impact AI Documentation: The Clock Expires

The 365-day window for agencies to document minimum risk management practices for high-impact AI closes April 3, 2026. OMB has not announced a formal public reporting mechanism for compliance outcomes. What agencies report internally, and how that information is used, will tell us whether the deadline produced a documentation event or a governance change.

PROCUREMENT

M-25-22 Acquisition Requirements Now in Effect

M-25-22's procurement requirements apply to contracts awarded on or after October 1, 2025, covering performance-based acquisition, ongoing AI system monitoring, and vendor disclosure of unanticipated AI use. Whether contracting officers are applying these provisions consistently is an open question. GSA's required procurement tools repository deadline has also passed.

NIST

AI RMF Revision in Progress: The Next Comment Window

The NIST AI RMF is in active revision. M-25-21 does not reference the AI RMF by name as required guidance, signaling a preference for flexibility over framework mandates. Whether the revised RMF is formally incorporated into future OMB guidance, or remains a voluntary reference, will shape federal AI governance infrastructure for years. The next public comment window is expected Q2-Q3 2026. GIAG is monitoring.

Government AI in Practice

Published by ThinkCapital LLC | thinkcapital.org | Michael Bragen, Principal

To participate in the GIAG research or share this newsletter, visit thinkcapital.org/research.html