Government AI in Practice

Substack Posts Archive · May 2026

Michael Bragen

Principal, ThinkCapital LLC

Government IT & AI Governance Initiative (GIAG)

thinkcapital.substack.com · thinkcapital.org

michael.bragen@thinkcapital.org

This archive collects posts from the *Government AI in Practice* newsletter and associated LinkedIn commentary, produced under the GIAG research initiative. The initiative examines how federal and state agencies implement AI governance frameworks, with particular focus on NIST AI RMF implementation fidelity (Stream One) and human oversight quality in agentic AI deployments (Stream Two).

19	4	2
Posts	Newsletter Issues	Research Streams

Compiled May 2026 · ThinkCapital LLC · Belmont, California

Contents

May 4, 2026 · *LinkedIn Post*
The M-25-21 Compliance Event: What It Measured and What It Did Not

May 5, 2026 · *LinkedIn Post*
The Identical Score Problem

May 6, 2026 · *Newsletter Article*
GQM Applied to Federal AI Governance

May 7, 2026 · *LinkedIn Post*
Early Stream One Patterns

May 8, 2026 · *LinkedIn Post*
NIST AI RMF vs. EU AI Act: Three Operational Differences

May 11, 2026 · *LinkedIn Post*
The CAO Designation and the ISSO Precedent

May 12, 2026 · *Newsletter Article*
The CAO Authority Gap

May 13, 2026 · *LinkedIn Post*
NASCIO's Decision Rights Standard

May 14, 2026 · *LinkedIn Post*
The FITARA Precedent for the CAO Authority Gap

May 15, 2026 · *LinkedIn Post*
Three Diagnostic Questions for Government AI Governance Leaders

May 18, 2026 · *LinkedIn Post*
The Intervention Point Problem

May 20, 2026 · *LinkedIn Post*
Two Oversight Models

May 20, 2026 · *Newsletter Article*
Issue 7: The Intervention Point Problem

May 21, 2026 · *LinkedIn Post*
Three Requirements for Oversight That Functions as a Genuine Control

May 22, 2026 · *LinkedIn Post*
Five Diagnostic Questions for Government AI Oversight

May 26, 2026 · *LinkedIn Post*
The NASCIO Paradox

May 27, 2026 · *Newsletter Article*
Issue 8: The NASCIO Paradox — 88% Have Frameworks, 75% Have Serious Concerns

May 28, 2026 · *LinkedIn Post*
The CPO Authority Gap in State AI Governance

May 29, 2026 · *LinkedIn Post*

Shadow AI and the Local Governance Gap

May 4, 2026 · LINKEDIN POST

The M-25-21 Compliance Event: What It Measured and What It Did Not

400 US federal agencies spent the past year complying with a new AI governance directive. Last month, the deadline passed. Did any of this make a difference?

The Office of Management and Budget Memorandum M-25-21 required every executive branch agency to inventory its AI use cases, designate a Chief AI Officer, update governance policies, and complete required compliance documentation by April 3, 2026. Those agencies filed their inventories, designated their CAOs, and updated their policies. By every process metric, the deadline was met. What nobody has yet answered: did any of it change a single AI deployment decision?

That is not a rhetorical question. It is the right question, and the compliance framework was not designed to ask it. M-25-21 established a deadline with documentation requirements and no verification architecture. No public compliance dashboard exists. No OIG audit program was announced with enforcement authority attached. No mechanism distinguishes an agency that used its CAO designation to change how deployments are approved from one that used it to satisfy a filing requirement.

This is the FISMA pattern. The Federal Information Security Management Act generated years of documentation activity and modest security improvement before measurement infrastructure caught up. The first FISMA evaluation reports said so explicitly. M-25-21 is FISMA applied to AI governance, one month in.

Wednesday's issue of Government AI in Practice examines what the compliance event actually produced, what governance documentation cannot measure, and what a functional measurement standard for federal AI governance would require.

If you are a government IT leader or CAO responsible for AI governance: has your governance program changed a production deployment decision in the past 12 months? If you cannot name a specific example, your program may be producing documentation rather than control.

May 5, 2026 · LINKEDIN POST

The Identical Score Problem

Consider two agencies.

Agency A completed every M-25-21 requirement: use case inventory filed, Chief Administrative Officer (CAO) designated, governance policy updated. In four months of operation, its benefits routing AI system has passed every scheduled review. The system has never been delayed, scoped down, or modified because of a governance review.

Agency B has the same documentation. It also has a CAO with defined decision authority over new AI deployments, a governance calendar that triggers operational review when production system behavior diverges from baseline, and a documented record of three deployment decisions that governance changed over the same period.

Under current M-25-21 measurement practice, these two agencies are indistinguishable. Both report full compliance. An auditor using current assessment criteria cannot tell them apart.

Agency A has governance outputs. Agency B has governance outcomes. The compliance architecture measures only the former.

This is not an implementation failure. The agencies that filed documentation accurately reported what they did. The framework was not designed to ask whether what they did changed decision behavior. It was designed to ask whether required activities were completed on schedule. Those are different questions with different answers.

Wednesday's issue of Government AI in Practice introduces the measurement discipline that would distinguish them, drawn from Victor Basili's Goal-Question-Metric framework, created to address this type of problem at NASA in the 1980s.

May 6, 2026 · NEWSLETTER ARTICLE

GQM Applied to Federal AI Governance

The measurement problem in federal AI governance has a 50-year-old solution. Victor Basili and colleagues at NASA developed the Goal-Question-Metric framework in the 1970s and 1980s to solve exactly the problem government AI governance now faces: organizations producing quality process outputs with no way to verify that those outputs reflected actual quality improvement.

GQM works backward from outcomes. Define the goal. Derive the questions that would confirm whether the goal is being met. Identify the metrics that answer those questions. When applied to AI governance, the goal is to ensure AI deployments operate within defined risk boundaries and that human accountability is preserved where it is required.

The questions that would confirm whether that goal is being met:

Has governance review changed a deployment decision in the past 12 months?

Is there a named individual with defined authority to halt or modify a production system outside the normal governance calendar?

When the system processes cases outside its training distribution, what triggers a review?

What proportion of escalation events result in a human override?

These questions have answers that distinguish agencies producing governance documentation from agencies producing governance outcomes. The Office of Management and Budget Memorandum M-25-21 asks none of them.

[Excerpt — full text available at thinkcapital.substack.com]

May 7, 2026 · LINKEDIN POST

Early Stream One Patterns

Stream One structured interviews are in progress. The findings here are early-stage, drawn from practitioner intake conversations and the public deployment record. These are directional patterns, not conclusions.

Four patterns are emerging across agencies.

Governance is front-loaded, not operational.

Risk assessment, authorization, and policy development receive substantial organizational attention before deployment. Post-deployment operational monitoring receives substantially less, and in many agencies is thin and unassigned. M-25-21 reinforced this: the deadline required documentation production, not monitoring design. The compliance event was itself front-loaded governance.

The CAO role carries compliance authority, not decision authority.

Most early-intake conversations describe a CAO responsible for ensuring required documentation is produced and procedures are documented. What CAOs typically do not hold is authority to delay or modify a deployment the program office wants to proceed. Compliance designation and operational authority are different things. Most CAO positions reflect the former.

Calendar triggers dominate behavioral triggers.

Governance reviews occur when the compliance calendar requires them. They rarely occur when the production system generates a signal that warrants review. A system encountering distributional shift — processing cases outside its training distribution — does not trigger a governance review in most deployments. The audit cycle drives oversight; system behavior does not.

Human-in-the-loop language remains undefined.

The term appears in M-25-21 and in virtually every agency governance document reviewed. It almost never specifies which decisions require human review, who conducts the review, what information the reviewer is given, or what the documentation standard is. It functions as a policy commitment. It does not create an operational control.

None of these patterns is a compliance failure. They are features of how the current framework is designed. The framework measures process completion. These patterns are what process-completion governance looks like in practice.

May 8, 2026 · LINKEDIN POST

NIST AI RMF vs. EU AI Act: Three Operational Differences

M-25-21 references the NIST AI RMF as the governance foundation for federal AI programs. Understanding what the RMF actually measures, and where it is silent, matters for any practitioner whose accountability extends beyond domestic compliance.

Three differences between the NIST AI RMF and the EU AI Act are operationally significant for US government agencies right now.

Oversight standard.

The RMF references oversight through the GOVERN and MANAGE functions. Neither specifies what oversight must accomplish. The EU AI Act requires technical measures that enable human oversight throughout the system lifecycle. That is a different standard, and it binds operators legally.

Accountability role.

M-25-21 mandates CAO designation. Neither M-25-21 nor the RMF specifies what operational authority the role must carry. The EU AI Act names operators with defined obligations, specifying what operators must be able to do, not just that a role must exist.

Verification.
Under the RMF, agencies self-assess. Documentation is the evidence of compliance. No external verification standard exists for framework claims. The EU AI Act requires post-market monitoring and incident reporting to a national authority. External verification exists; documentation alone is not sufficient.

For US agencies with EU-facing deployments or vendor relationships that cross jurisdictions, the compliance gap is not theoretical. It is operational now.

The practitioner implication: an agency can satisfy every RMF and M-25-21 requirement with no evidence that governance changed a deployment decision. The framework does not require that evidence. The EU AI Act does.

May 11, 2026 · LINKEDIN POST

The CAO Designation and the ISSO Precedent

When the Federal Information Security Management Act (FISMA) was implemented, agencies designated Information System Security Officers as the named accountability point for security at the system level. The ISSO became one of the most reliably documented positions in federal IT governance. It also became, in many agencies, a compliance position: the person who coordinates audits, produces required documentation, and ensures the security checklist stays current.

Whether the ISSO held authority to halt a production system whose security posture had degraded below the authorization baseline was a different question, rarely answered the same way twice.

M-25-21's Chief AI Officer (CAO) requirement is replicating this architecture at speed.

The designation requirement is clear. The authority standard is absent. A CAO who can require documentation but cannot delay a deployment the program office wants to proceed is an ISSO with an AI-specific title. The compliance record looks identical to a CAO with real deployment authority.

Stream One research is designed to distinguish them. Tomorrow's issue of Government AI in Practice examines what the M-25-21 designation actually requires, what NASCIO's state CIO operating model reveals about decision rights, and what the public record of CAO job postings confirms about where organizations have drawn the line.

May 12, 2026 · NEWSLETTER ARTICLE

The CAO Authority Gap

When your agency's first significant AI governance failure occurs, accountability will be allocated based on what your Chief AI Officer (CAO) was authorized to do, not what the CAO was designated to oversee.

The program office will show the CAO reviewed the documentation. The CAO will show concerns were raised. No one will show the system was halted.

That is a paper trail, not a governance outcome.

M-25-21 required federal agencies to designate a CAO. Most did. The directive specified who must be named. It did not specify what that person can stop. An agency with a compliance-only CAO and an agency with a CAO holding real deployment decision authority produce the same M-25-21 compliance record. The designation looks identical. The accountability will not be.

[Excerpt — full text available at thinkcapital.substack.com]

May 13, 2026 · LINKEDIN POST

NASCIO's Decision Rights Standard

The National Association of State Chief Information Officers (NASCIO) represents the state technology executives responsible for IT policy and strategy across all 50 states. Its research sets the operational and governance standards that state CIO offices use as reference architecture. When NASCIO publishes a framework, state IT leaders read it.

NASCIO's "Evolving Role of the State CIO as Change Leader" report, published April 21, goes beyond what the OMB memorandum M-25-21 dictates: an explicit definition of decision rights attached to a governance role.

The NASCIO model defines what the state CIO owns across three operating layers. The Executive Integration Layer names a specific right: "Authorizes movement between exploration and production." That is the AI governance equivalent of approving a deployment decision. NASCIO names it as belonging to the CIO. Contrast that to the OMB memo which names no equivalent right for the CAO.

The NASCIO framework provides a concrete benchmark: a governance role defined with explicit authority over the transition from pilot to production. Stream One is using that standard to evaluate CAO authority across federal and state agencies. Early-intake interviews are confirming the gap.

May 14, 2026 · LINKEDIN POST

The FITARA Precedent for the CAO Authority Gap

For government Chief Information Officers (CIOs) and Chief AI Officer (CAO) designees, OMB Memorandum M-25-21 compliance is not the same as governance authority. The agencies best positioned when the first significant AI governance failure occurs will be those that have explicitly defined what their CAO is authorized to stop, not just required to document. That definition is an organizational decision. M-25-21 does not make it for you.

Federal IT has a direct precedent for what happens when a governance role is designated without specified authority. The Federal Information Technology Acquisition Reform Act (FITARA), enacted in 2014, was designed to give federal CIOs the authority to direct and oversee their agencies' IT acquisitions. Legislation was necessary because the CIO designation had existed for years without the operational authority the role nominally required. Agencies had CIOs. The authority to make consequential technology decisions remained distributed across program offices.

The Government Accountability Office (GAO) documented what followed. Four years after FITARA's enactment, GAO found that none of the 24 agencies reviewed had established policies fully addressing the role of their CIO. None of the 27 resulting recommendations had been implemented as of June 2019. GAO's 2025 High-Risk Series confirmed the underlying problem: FITARA was intended to strengthen the authority of CIOs to provide needed direction and oversight of covered agencies' IT acquisitions. Closing the gap required legislation and, even then, took years to produce documented results.

M-25-21 created the CAO designation in the same architecture. The directive names who must be designated. It does not specify what deployment decisions the CAO may delay or halt, what authority the CAO holds relative to program offices, or what recourse exists when a CAO judgment is overridden. Left unaddressed, the FITARA record predicts agency-level discretion will produce widely inconsistent authority structures, the compliance record will not surface the difference, and accountability questions will arrive through post-incident review rather than governance design.

The compliance record will not distinguish an agency that built a CAO role with real operational authority from one that did not. The post-incident review will. Define the authority structure now, before you need to defend its absence.

May 15, 2026 · LINKEDIN POST

Three Diagnostic Questions for Government AI Governance Leaders

Three questions every government CIO, CAO, and AI governance leader should be able to answer.

Question 1: What is your CAO's authority when a governance review produces a finding the program office disputes?

If the CAO documents the finding and the program office decides whether to act, you have compliance authority. If the CAO can require a pause while the finding is resolved, you have coordination authority. If the CAO can halt a deployment pending governance resolution, you have operational decision rights. M-25-21 requires none of these specifically.

Question 2: Has your CAO delayed or modified a deployment decision in the past 12 months?

Not recommended a modification. Delayed or modified it. A CAO who has never changed a deployment decision may hold governance responsibility without governance authority. The compliance record does not distinguish those two situations.

Question 3: Does your CAO hold the right to authorize movement from AI pilot to production deployment?

NASCIO defines this as a named CIO decision right. M-25-21 does not specify whether the CAO holds an equivalent. If your agency has not assigned this authority explicitly, program offices hold it by default.

These questions are drawn from the Stream One research design. They are what Stream One is building the first comparative empirical record of.

May 18, 2026 · LINKEDIN POST

The Intervention Point Problem

Government AI governance frameworks require human oversight. Very few of them define what that oversight is supposed to do.

That gap has a specific consequence.

When the first significant failure occurs in a government agentic AI deployment, investigators will find documentation showing a review process existed. The harder question will be whether the reviewer was positioned to stop what went wrong, or whether they were looking at outputs after the decisions that caused the failure had already been made.

In agentic AI systems, consequential decisions happen inside process chains, not at output boundaries. A system executing connected actions across multiple data sources makes choices at every step. Reviewing the output tells you whether the final result looks right. It does not tell you whether the process that produced it stayed within appropriate limits.

Most government AI oversight frameworks position review at the end of that process. Some do not specify where review should be positioned at all.

GIAG Stream Two research is examining where agencies actually place human oversight relative to where consequential decisions occur. The pattern across every early-intake interview conducted to date: review is downstream from the decisions.

May 20, 2026 · LINKEDIN POST

Two Oversight Models

Every government agency deploying AI has adopted — either deliberately or by default — one of two oversight models.

Model one: a gate before deployment and a log review after. The system runs between those two endpoints without structured human engagement. The approach is to authorize before and audit after.

Model two: human review gates are positioned at consequential decision nodes inside the process, before the system continues to the next step. Human authorization intervenes at the point where choices are being made.

Empirical findings indicate that more than half of agencies use the first model, and an increasing number have governance documents that describe model two.

When your organization places oversight at the edges, the risk of scope drift increases dramatically. A government AI system authorized at deployment with a defined scope may, six months later, have more integrations, data sources, and task delegation. Each addition looks minor in isolation and does not trigger formal re-authorization. In aggregate, they produce a system doing significantly more than what was approved.

Drift is avoided when oversight mechanisms are calibrated against the original scope.

Issue 7 documents four cases that show the same structural pattern: the US Customs and Border Protection Automated Targeting System, the Transportation Safety Administration's facial recognition program, the State of Arkansas's Medicaid care algorithm, and the Veterans Administration claims processing AI. In every case, scope expanded but oversight did not.

The question for any government IT leader: is your agency tracking the divergence between authorized scope and actual operating scope for each deployed system? If not, someone is reviewing against the wrong specification.

May 20, 2026 · NEWSLETTER ARTICLE

Issue 7: The Intervention Point Problem

Issue 7 of Government AI in Practice examines what it means for human oversight to be positioned correctly in agentic AI deployments. The comparative oversight model framework, four documented scope drift cases, a set of practitioner diagnostic questions, and three requirements for oversight that functions as a genuine control are in the full issue.

[Excerpt — full text available at thinkcapital.substack.com]

May 21, 2026 · LINKEDIN POST

Three Requirements for Oversight That Functions as a Genuine Control

Most government agentic AI deployments have at least one of three diagnosable gaps. Many have all three: a governance gap (unclear accountability, oversight, or intervention authority); an operational gap (weak integration into operational workflows and decision processes); and an adoption gap (low organizational readiness, trust, or behavioral alignment).

For governance to function as real operational control rather than simply documenting assumptions, these gaps must be addressed directly.

Intervention points inside the process chain.

Reviews must occur at the specific decision nodes where the system makes choices that shape downstream outcomes. Process-level errors are often invisible to output-only review. If oversight occurs only after an output is produced, the critical decisions that created the outcome have already taken place.

An operationally defined scope, actively monitored.

Deployment scope is not a snapshot. As integrations accumulate and task delegation expands, the authorized scope continuously shifts. Without ongoing monitoring, reviewers may evaluate the wrong specification. Most frameworks assess scope periodically rather than continuously, and often only after an audit or incident.

Reviewers equipped for the actual task.

Reviewing an agentic AI process for integrity and scope adherence is a different cognitive job than reviewing a classification output. The information required — and the expertise required — is different. Naming a reviewer without specifying what they are evaluating and what authority they hold to act on a

finding is not oversight. It is a record of presence.

These gaps are often visible early. In many cases, a focused assessment is enough to clarify whether the challenge is technical, organizational, or a mismatch between ambition and readiness. Organizations that recognize these conditions early are in a much stronger position to move beyond isolated experimentation toward sustainable operational use of AI.

May 22, 2026 · LINKEDIN POST

Five Diagnostic Questions for Government AI Oversight

Five questions worth sitting with before Monday. These come directly from GIAG Stream Two research conversations with government technology practitioners responsible for AI oversight.

One.

Can you identify the specific points in your deployed AI workflows where a human decision is required before the system continues execution?

Two.

Has the operating scope of your current AI deployments changed since initial authorization? If so, has that change been formally reviewed?

Three.

When your reviewers evaluate AI system outputs, do they have visibility into the process steps that produced them, or are they working from results only?

Four.

Has your agency defined what a reviewer does when uncertain about an AI-generated output? A documented protocol with binding authority, not a general permission to escalate.

Five.

How does your governance framework distinguish between oversight of assistive AI and oversight of agentic AI? Or does it apply a single oversight model to both?

These questions locate where oversight architecture is thinner than governance documentation suggests.

May 26, 2026 · LINKEDIN POST

The NASCIO Paradox

NASCIO's 2025 State CIO Survey contains a number that deserves more attention than it has received.

88% of states report having AI responsible use policies in place. 75% of state CIOs report serious concerns about deploying GenAI in direct citizen services. Both numbers come from the same survey population.

If the governance frameworks were resolving the risks they were designed to manage, the serious concerns should be diminishing. They are not. The documentation layer is present, but the confidence layer does not follow from it.

Tomorrow's issue of Government AI in Practice examines what accounts for that gap at the federal, state, and local level, and what closing it requires.

May 27, 2026 · NEWSLETTER ARTICLE

Issue 8: The NASCIO Paradox — 88% Have Frameworks, 75% Have Serious Concerns

88% of states have AI responsible use policies. 75% of state CIOs have serious concerns about deploying GenAI in citizen services. Those numbers are consistent with each other. High policy coverage and low deployment confidence describe the same condition: documentation without operational architecture.

Issue 8 examines what that gap looks like at each tier of government and what it takes to close it.

Federal.

What subnational governance frameworks anchor to when the federal signal goes quiet. And the one legislative anchor that remains stable.

State.

Why CPO authority, inventory currency, and undefined review roles are where state governance frameworks break down in practice.

Local.

How documentation-only frameworks cascade downstream. Why shadow AI hits hardest where governance capacity is thinnest.

[Full issue at thinkcapital.substack.com]

May 28, 2026 · LINKEDIN POST

The CPO Authority Gap in State AI Governance

31 states have a Chief Privacy Officer playing a central role in AI governance. The NASCIO April 2026 CPO survey finds that most of those CPOs lack dedicated funding, defined authority over AI deployment decisions, and sufficient staffing to perform the oversight function their governance documents assign them.

The practical implication: naming a role in a governance document is a starting point, not a completion. Three questions are worth asking about any named oversight role in your framework right now.

Can the CPO act, or only advise?

Advising on a recommendation engine and intervening in a deployed agentic system that takes autonomous multi-step action are different functions. Most governance frameworks have not drawn that line, and most CPO authority structures were not built with agentic AI in mind.

Does the CPO have budget and staffing proportionate to the oversight scope?

Dedicated funding and staffing are not administrative details. They determine whether oversight is operational or nominal.

What specifically can the CPO do when they find a problem mid-deployment?

The answer to that question should be in writing, with defined criteria, access rights, and a timeframe for action.

Issue 8 examines the CPO authority gap alongside two related structural problems: inventory currency and undefined reviewer authority. All three are measurable. None of them are resolved by adding policy language.

May 29, 2026 · LINKEDIN POST

Shadow AI and the Local Governance Gap

Local government agencies typically have the least AI governance capacity and the fewest dedicated oversight resources. They are also where informal AI adoption most consistently outruns formal governance. CompTIA research released this week identifies unauthorized AI tool use in government as a growing risk. The exposure is not evenly distributed.

This week's HHS announcement sharpens the problem. The Administration for Children and Families is offering $6 million to state and local governments to upgrade predictive analytics capabilities in child welfare systems — federal funding accelerating algorithmic decision-making in high-stakes services for vulnerable populations, with no apparent governance requirements or bias mitigation mandates attached. That is the governance capacity gap made concrete.

Local agencies are downstream from state governance frameworks. They adapt state policy rather than build from first principles. When state frameworks include operational architecture — funded oversight roles, current inventories, defined reviewer authority — local agencies inherit a functional floor. When state frameworks consist primarily of documentation, local agencies inherit the documentation. The operational protection the documentation describes is not part of what transfers.

NIST AI RMF 1.0 is the most practical available baseline for local agencies that need governance infrastructure independent of state or federal signal uncertainty. The value is in using it as an operational foundation, not a compliance checklist. Agencies that treat it as a checklist produce more documentation. Agencies that treat it as architecture produce governance that functions regardless of what the tiers above them are doing.

The CBS News poll finding that most Americans do not trust government to use AI appropriately reflects the aggregate output of a system where governance documentation has outrun governance architecture at every level. Agencies that close that gap are the ones that will change that number.

About the GIAG Initiative

The Government IT and AI Governance Initiative (GIAG) is an independent research program operated under ThinkCapital LLC, examining how federal and state agencies implement artificial intelligence governance frameworks in practice.

Stream One investigates NIST AI RMF implementation fidelity — whether documented governance programs produce operational governance outcomes, and what structural factors determine the gap between the two.

Stream Two investigates human oversight quality in agentic AI deployments — how agencies define, assign, and sustain meaningful human control as AI systems take on expanded operational roles.

The initiative publishes the *Government AI in Practice* newsletter on Substack and maintains working papers and research documentation at thinkcapital.org.

Michael Bragen | Principal, ThinkCapital LLC
michael.bragen@thinkcapital.org
thinkcapital.org · thinkcapital.substack.com
www.linkedin.com/in/bragen

Belmont, California · Compiled May 2026