GOVERNMENT AI IN PRACTICE

Research and analysis from the ThinkCapital GIAG Initiative

ISSUE 6 · MAY 12, 2026

EARLY SIGNAL: FROM THE RESEARCH

Office of Management and Budget (OMB) Memorandum M-25-21 required federal agencies to designate a Chief AI Officer (CAO) by April 3, 2026. Most did. The National Association of State Chief Information Officers (NASCIO) published its Change Leader report on April 21 documenting state governments building parallel CAO structures. What neither M-25-21 nor the NASCIO framework specifies is *what operational authority the CAO role must carry*. Stream One practitioners describe a role that holds compliance responsibility without deployment decision authority. The designation is present, but its function is undefined.

Pattern 2 from Stream One early-stage intake is holding as interviews continue: no practitioner interviewed to date has described a case in which the CAO delayed or modified a deployment over program office objection. The authority standard was not specified in M-25-21. Most agencies defaulted to the compliance model. That is the rational response to a compliance architecture that does not reward the stronger model.

FROM THE EDITOR

The CIO reading this faces a decision that looks administrative but is entirely about strategy. M-25-21 requires every covered agency to designate a Chief AI Officer (CAO). How you define that CAO's authority right now is how accountability will be allocated after the first significant failure. A CAO who can flag but not stop is a witness, not a safeguard. Most agencies have not yet had a public AI governance failure. That window, after designation and before crisis, is when governance architecture gets built on principle rather than in response to pressure. It is closing.

The CAO designation arrived with a clear mandate and an absent authority standard. The regulation indicates who must be named, but does not specify what that person can stop. Federal IT has been here before.

When FISMA was implemented, agencies designated Information System Security Officers (ISSOs) as the named accountability point for security at the system level. The role became one of the most reliably documented positions in federal IT governance, and in many agencies, a compliance position: the person who coordinates audits, produces required documentation, and ensures the checklist stays current. Whether the ISSO held authority to halt a production system whose security posture had degraded was a different question, rarely answered the same way twice. The CAO designation is replicating this architecture. A CAO who can require documentation but cannot delay a deployment is an ISSO with an AI-specific title.

The accountability stakes are not theoretical. A clean ISSO record did not protect agency leadership when breaches occurred. It sharpened the question: if every box was checked, why did this happen? The compliance record became evidence of a deeper problem. The CAO designation will work the same way. 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 merely a paper trail, not a governance outcome.

This issue examines what the designation requires, what NASCIO's state CIO model reveals about decision rights, and what CAO job postings confirm about where agencies have drawn the line.

~ Michael

WHAT M-25-21 REQUIRES OF THE CAO

The CAO provisions in M-25-21 are worth examining precisely. The directive requires each covered agency to designate a Chief AI Officer responsible for coordinating AI governance activities, ensuring AI use cases are inventoried, and serving as the agency's primary accountability point for AI governance compliance. The designation requirement is specific. The authority standard is absent.

M-25-21 does not specify what decisions the CAO may delay or halt, what authority the CAO holds relative to program offices, whether the CAO requires concurrence before a deployment proceeds, what the CAO's authority is when a governance review conflicts with a program office timeline, or what recourse exists when a CAO judgment is overridden. The role is defined entirely by its compliance obligations. Operational authority is left to agency discretion.

This is not an oversight, it’s a conscious choice. OMB directives that specify role designations without authority standards have a consistent precedent in federal IT governance. The Federal chief information officer position was established with portfolio coordination responsibility and limited operational authority over agency IT decisions until the Federal Information Technology Acquisition Reform Act (FITARA) strengthened those provisions in 2014. The authority gap persisted for more than a decade. M-25-21 created the CAO designation in the same architecture.

THE NASCIO OPERATING MODEL: DEFINED DECISION RIGHTS

The NASCIO “Evolving Role of the State CIO as Change Leader” report, published April 21, 2026, describes a three-layer operating model for state CIOs. Understanding the model is useful for the CAO question not because it describes authority tiers, but because it contains something M-25-21 does not: an explicit definition of decision rights attached to a governance role.

The NASCIO model describes three layers:

- The Core Operations Layer handles stability, efficiency, compliance, infrastructure, and service delivery.
- The Exploration and Modernization Layer covers digital innovation, AI pilot programs, and modernization experiments.
- The Executive Integration Layer addresses portfolio governance, funding, and strategic decisions about which pilots advance to production.

The report explicitly defines that each layer has distinct accountability: “Each horizon has distinct resource mechanisms, metrics and decision rights.” (NASCIO, April 2026, p.7)

The report's Decision Rights and Accountability section defines what the state CIO owns. The listed CIO decision rights include:

- Owning enterprise technology strategy and standards

- Ensuring legal, security, and ethical alignment
- “Authorizes movement between exploration and production.” (NASCIO, April 2026, p.9) This is the AI governance equivalent of approving a deployment decision. NASCIO names it as belonging to the CIO.

M-25-21 names no equivalent decision right for the CAO. The directive designates the role. It does not specify that the CAO holds authority to authorize, delay, or halt a production AI deployment. The NASCIO model shows what it looks like when a governance role has defined decision rights. M-25-21 shows what it looks like when it does not. That contrast is directly relevant to what Stream One is finding in practice.

THE FISMA ISSO ANALOGY

The ISSO role emerged from FISMA's requirement for a named accountability point at the system level. The ISSO maintains the system's authorization documentation, coordinates required reviews, manages Office of Inspector General (OIG) interactions, and ensures the security posture is documented against the baseline. These are compliance functions. They produce a reliable documentation record.

What the ISSO role does not standardly carry is authority to halt a production system whose security posture has degraded below the authorization baseline. That authority gap produced a consistent pattern: a reliable documentation record and inconsistent translation of security findings into actual deployment decisions. The Government Accountability Office (GAO) documented this pattern repeatedly across FISMA review cycles (through named security officials, clear compliance responsibility, and an operational boundary that agencies interpreted inconsistently.)

THE GAO FITARA RECORD: DESIGNATION WITHOUT AUTHORITY

GAO documented the federal CIO authority gap over more than a decade of FITARA review. Two findings from that record are directly relevant to the CAO question:

First, on the scope of the problem: GAO reported in August 2018 (four years after FITARA was enacted to address CIO authority) that none of the 24 selected agencies had established policies that fully addressed the role of their CIO. None of the 27 resulting GAO recommendations had been implemented as of June 2019. Designation had not produced authority. [Source: CRS, Federal Information Technology Acquisition Reform and Management, R44843]

Second, on why FITARA was necessary: GAO's 2025 High-Risk Series report states that “FITARA was intended to strengthen the authority of CIOs to provide needed direction and oversight of covered agencies' IT acquisitions.” [GAO-25-107852] The word “strengthen” confirms that the pre-FITARA CIO lacked the authority the role nominally required. The designation existed. The authority required legislation.

M-25-21's CAO designation is at the pre-FITARA stage of this pattern. The role has been named, but Its authority has not been legislated or otherwise specified. The FITARA precedent predicts that the gap will require deliberate action to close but does not predict how long this will take.

M-25-21's CAO structure is replicating this architecture. The CAO is a named compliance accountability point. The designation requirement is clear. The operational boundary is undefined. A CAO who receives a finding from a governance review and cannot act on it without program office concurrence is exercising ISSO-level authority. The documentation looks the same as a CAO with deployment decision rights. The compliance record cannot tell them apart.

WHAT PUBLIC CAO JOB POSTINGS REVEAL

A review of public Chief AI Officer job postings from federal and state agencies published between January and April 2026 shows a consistent pattern. Listed responsibilities concentrate in three areas: governance documentation and compliance management, stakeholder coordination and policy communication, and AI use case inventory development and oversight. Named authority over deployment decisions appears in a minority of postings and is described in advisory terms when it does.

The most common authority language in postings reviewed is coordination-layer language: the CAO will work with program offices, ensure required governance activities occur, and represent the agency in interagency AI governance forums. Decision-layer language specifying that the CAO holds authority to delay, modify, or halt a deployment appears rarely. Several postings describe the CAO's governance role in terms identical to the ISSO model: responsible for ensuring required documentation is produced and that the agency meets applicable compliance standards.

The postings reveal where agencies have drawn the line as a matter of organizational design, not just regulatory interpretation. An agency that posts a CAO role with coordination authority has made an organizational decision about how the role functions. M-25-21 compliance is satisfied either way. The governance effectiveness will be significantly different depending on what choice is made.

AUTHORITY STRUCTURE COMPARISON: CAO VS. ISSO VS. FEDERAL CIO

The table below maps three federal IT governance roles against five dimensions. The purpose is to make explicit where the M-25-21 CAO designation sits relative to roles where the authority question has been formally addressed.

	Authority Dimension	Federal CIO (post-FITARA)	CAO (M-25-21)
1	Compliance accountability	Named and defined by statute.	Named and defined by directive.
2	Deployment decision authority	Defined for major IT investments by FITARA; agency-variable elsewhere.	Undefined. Left to agency discretion.
3	Authority to delay a program office deployment	Present in FITARA domains; contested outside them.	Not specified. No standard established.
4	Authority standard in governing document	Strengthened by FITARA in 2014; GAO found gaps persisted for years after enactment.	Absent. Designation requirement only.
5	Stream One finding	Not directly studied; pre-FITARA CIO gap is analogous.	CAOs report compliance authority. Deployment authority is rare and informal.

THE IDENTICAL SCORE PROBLEM, APPLIED TO THE CAO ROLE

An agency with a CAO holding compliance-only authority and an agency with a CAO holding deployment decision authority produce identical M-25-21 compliance records. Both have a designated

CAO. Both can report that designation in required documentation. The compliance architecture distinguishes neither.

The measurement problem compounds the governance problem. An agency that has invested in building a CAO role with real deployment decision authority has no way to signal that investment in the current compliance record. Its CAO and a compliance-only CAO at a peer agency look identical in every official reporting channel. For government CIOs and agency heads who want their governance programs to go beyond compliance and to be credible, that indistinguishability is a direct cost. It removes the incentive to build the stronger model.

Stream One research is building the first comparative empirical record of where CAO authority actually sits across agency types. The findings will be published in Working Paper WP-3, currently in development. The central question is not which agencies designated a CAO. It is which CAOs changed a deployment decision.

STREAM ONE: PATTERN 2 UPDATE

Issue 5 identified Pattern 2 from early-stage Stream One interviews: the CAO role in most early-intake conversations holds compliance responsibility without deployment decision authority. The pattern is holding as structured interviews continue.

The consistent description across agency types is a CAO who coordinates documentation, manages the use case inventory, and serves as the point of contact for governance audit activity. No early-intake practitioner has described a case in which the CAO halted or modified a deployment over program office objection. Several have described cases in which the CAO raised a governance concern and the program office timeline was not changed. That is advisory authority, which does not equate to operational control.

The pattern distinction matters because it is not a criticism of individual CAOs or agencies. It is a structural observation about what M-25-21's designation requirement created. The authority standard was not specified. Most agencies defaulted to the compliance model. That is the rational response to a compliance architecture that does not reward the stronger model.

PRACTITIONER DIAGNOSTIC: THREE QUESTIONS FOR GOVERNMENT IT LEADERS

The following questions are drawn from the Stream One research design and from practitioner intake conversations. They function as a rapid self-assessment for government CIOs, CAOs, and governance leadership responsible for AI programs.

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 a compliance-authority CAO. 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 does not require any specific level. Accountability standards eventually will.

2	**Has your CAO delayed or modified a deployment decision in the past 12 months?**	Not recommended modification. Not raised a concern. Delayed or modified it. A CAO who has never changed a deployment decision may have governance responsibility without governance authority. Or there may have been no findings that warranted intervention. The question distinguishes those two situations. The compliance record does not.
3	**Using the NASCIO Decision Rights framework as a reference: does your CAO hold the right to authorize movement from AI pilot to production deployment?**	NASCIO defines this as a named CIO decision right: "Authorizes movement between exploration and production." M-25-21 does not specify whether the CAO holds an equivalent right. If your agency has not explicitly assigned this authority to the CAO, the default is that program offices hold it. That is the most common pattern Stream One is finding.

STREAM TWO WORKING PAPER

The five-characteristic framework for human oversight in agentic AI deployment, as previously described includes irreversibility, consequence transfer, distributional novelty, value conflict, and legal or regulatory significance. For a full description and analysis of this framework, see the working paper, *When Humans Must Intervene: A Decision-Grounded Framework for Human Oversight in Government and Commercial Agentic AI Deployments*, available at thinkcapital.org/publications.html.

The CAO authority question maps directly to Stream Two research: a CAO without operational decision rights cannot exercise the intervention authority the five-characteristic framework identifies as structurally required for high-stakes agentic deployments. Compliance designation and intervention capability are different organizational constructs.

OUR PARTICIPATION ASK

Both GIAG research streams are actively recruiting government and public sector technology practitioners for structured research interviews. Stream One examines National Institute of Standards and Technology (NIST) AI Risk Management Framework (AI RMF) implementation fidelity and M-25-21 compliance outcomes across agency types. Stream Two examines human oversight quality in agentic AI deployments.

If you hold or have held responsibility for AI governance at a federal, state, or local agency, your perspective on the CAO authority question is directly relevant to this research. The structured interview is thirty minutes, conducted under a research confidentiality protocol. No agency or individual is identified in published findings.

Government AI in Practice is a practitioner research letter published by the Government IT and AI Governance Initiative (GIAG), a research program of ThinkCapital LLC. GIAG examines AI governance implementation in government and public sector contexts.

Michael Bragen, Principal, ThinkCapital LLC | michael.bragen@thinkcapital.org | thinkcapital.org | thinkcapital.substack.com