


# GOVERNMENT AI IN PRACTICE

Research and analysis from the ThinkCapital GIAG Initiative

ISSUE 8 · MAY 27, 2026



## EARLY SIGNAL: FROM THE RESEARCH

NASCIO's 2025 State CIO Survey puts two numbers in direct tension. Eighty-eight percent of states report they have AI responsible use policies in place. Seventy-five percent of state CIOs report serious concerns about deploying GenAI in direct citizen services. Both numbers come from the same survey population.

Taken together, these numbers reveal a gap that most governance discussions have not directly addressed. If the concerns remain serious, the frameworks are not resolving the risks they were designed to manage. Stream One research has been examining exactly this gap since January. This week's federal governance news adds a dimension to that examination.

## From the Editor

*Federal AI governance policy shifted again this week. A proposed executive order that would have established federal review authority for AI models before public deployment was withdrawn after industry objections that the provision would disadvantage U.S. developers in the international AI competition. There is now no federal pre-deployment review mechanism and no announced timeline for one.*

*State practitioners will read this alongside a CBS News poll released the same week showing that majorities across age groups and education levels do not believe government will ensure AI is used appropriately. These two data points are related. The federal governance signal is in flux. Public confidence in government AI stewardship is low. State and local agencies sit in the middle.*

*This issue examines what that condition means at three levels. The federal section is brief because the governance story there is, for now, one of absence rather than action. The state section gets most of the space, because that is where the more precisely documentable problem is. NASCIO's own survey data reveals a paradox that most state governance discussions have not fully confronted. The local section examines what state-level governance uncertainty produces downstream.*

~ Michael

## GOVERNANCE WITHOUT A FIXED ANCHOR

*What Federal Uncertainty, the NASCIO Paradox, and Downstream Risk Mean for Government AI at Every Level*

---

### FEDERAL

#### The Absent Signal and What It Exposes

Most state and local AI governance frameworks were built with federal standards as an implicit anchor. OMB guidance, NIST frameworks, and executive directives establish the vocabulary. State CIOs translate those requirements into local policy. Local agencies translate state policy into operational procedure. The translation chain depends on a federal signal.

When that signal is inconsistent or absent, governance frameworks calibrated to federal standards do not automatically recalibrate. They continue operating against a reference that may no longer reflect current federal direction. This is not a criticism of any particular policy decision. It is a structural observation about how governance architecture at subnational levels works in practice.

What this week's news does clarify: agencies whose governance frameworks consist primarily of compliance documentation referenced to federal standards are more exposed during periods of federal policy uncertainty than agencies that have built governance architecture that functions independently. The distinction between those two conditions is what Stream One research has been measuring since January.

One federal development worth noting for its contrast: the SBA AI accountability legislation passed Congress unanimously this year and establishes mandatory AI inventories, use-case reporting, and transparency requirements. The bipartisan support for AI accountability through legislation signals something distinct from the executive branch governance posture. CIOs who frame governance investment as alignment with the legislative direction rather than the current executive posture are on defensible ground. This legislative signal is one of the few stable federal anchors that state CIOs can reference.

---

### STATE

#### 75% Have Concerns, 88% Have Frameworks: The Paradox the Data Reveals

NASCIO's 2025 State CIO Survey reports that 88% of states have AI responsible use policies and 84% are actively inventorying AI uses. By every structural measure, state AI governance is in place. Policy documents exist. Inventories are underway. Responsible use language is present across governance frameworks.

The same survey reports that 75% of state CIOs have serious concerns about deploying GenAI in direct citizen services.

Those two numbers do not resolve each other. Governance frameworks and serious operational concerns coexisting in the same agency population is precisely what you would observe if the frameworks are not resolving the risks they were designed to manage. The documentation layer is present, but the confidence layer does not follow from it.

#### Why the Gap Exists

Stream One research has examined this pattern across federal, state, and local contexts. The consistent finding is that agencies produce governance documentation aligned with external frameworks without redesigning the oversight architectures those documents describe. A responsible use policy and an inventory exist; however, a mechanism that would ensure the policy is operative during system execution often does not. This absence is the intervention-point gap. It is a missing mechanism that would require scope monitoring by a reviewer with defined authority.

NASCIO's April 2026 CPO survey adds specific structural detail. Thirty-one states have a Chief Privacy Officer playing a central role in AI governance. Most of those CPOs lack dedicated funding, defined authority over AI deployment decisions, and sufficient staffing to perform the oversight function the governance

documents assign them. The governance document names the role, but the operational environment does not support it.

This is the same pattern the Four Cases cited in Issue #7 documented at the federal level: CBP's Automated Targeting System, TSA facial recognition, the Arkansas Medicaid algorithm, and VA claims processing AI. In each case, the governance documentation described oversight. The operational architecture did not sustain it. The pattern is not specific to any level of government. It is specific to governance built at the documentation layer without corresponding investment in the architecture layer.

### What Closing the Gap Requires

For state CIOs, the NASCIO data points to three specific investments that existing frameworks are not consistently driving.

**Operationally funded CPO authority.** A CPO who is named in a governance document but lacks dedicated budget, staffing, and clear authority to stop or modify deployments is performing a function that is more documentation than governance. The 31-state CPO figure means something different depending on whether those CPOs can act, or only advise. Governance documents that assign oversight roles without assigning the operational capacity to perform them are merely describing what oversight would look like if it were adequately resourced.

**Active inventory maintenance, not one-time cataloging.** Eighty-four percent of states are inventorying AI uses. The governance value of an inventory depends entirely on whether it tracks current operational scope or authorized scope at deployment. A system that was inventoried accurately in 2024 and has since added integrations, expanded data access, and taken on additional task delegation is no longer described by that inventory. The open question is whether an inventory is current.

**Specific definitions of reviewer authority.** Stream Two research has not yet found a single agency that has operationally defined what a reviewer does when they find a problem mid-process in a deployed agentic system. What information do they have access to? What specifically are they evaluating? What can they do with a finding, and within what timeframe? Responsible use policies that name a review role without answering those questions are inadequate for oversight.

---

## LOCAL

### Downstream From a Governance Gap

Generally, local agencies do not build AI governance frameworks from first principles. They adapt state frameworks, reference state guidance, and use state-procured systems. When state governance frameworks are strong, local governance inherits at least a partial floor. When state frameworks are primarily documentation without operational architecture, local agencies inherit the documentation without the protection it implies.

Shadow AI accelerates this problem at the local level. CompTIA research released this week identifies unauthorized AI tool use in government as a growing risk. Local agencies, which typically have the least governance capacity and the fewest dedicated AI oversight resources, are also the most likely to see informal AI adoption outrun formal governance. The combination of inherited documentation-only frameworks from the state level, inadequate local governance capacity, and accelerating informal AI adoption produces the highest governance-to-architecture gap in the system.

NIST AI RMF 1.0 is the most practical available option for local agencies that need a governance baseline independent of both state and federal signal inconsistency. It is not a complete answer to the intervention-point and scope-monitoring problems Stream Two research has identified. However, it provides a structured vocabulary that supports governance investment that does not depend on any particular state or federal mandate being active. Local CIOs using NIST AI RMF as an operational baseline rather than a compliance checklist are building governance that can survive policy uncertainty at the tiers above them.

The recent CBS News poll finding that most Americans do not trust government to use AI appropriately reflects aggregate public assessment of the full system. It is not a specific judgment about any tier of government. It is the observed output of a system where governance documentation has outrun governance architecture at every level simultaneously. That is the condition the research is documenting. It is also the

condition that creates the opportunity: agencies that close the gap between their governance documentation and their governance architecture are the ones that will gain the public's trust.

### APPLIED RESEARCH: THE GIAG ASSESSMENT TOOL AND WP3

GIAG Working Paper 3 (Mandate Translation) releases in early June. The paper examines how federal AI governance requirements arrive at state and local agencies and what happens to their operational substance in the translation process. The central finding: most agencies produce governance documentation aligned with federal standards while leaving oversight architectures operationally unchanged. This issue's state section is a preview of that argument applied to the NASCIO data.

The GIAG Assessment Tool at [thinkcapital.org/tools.html](https://www.thinkcapital.org/tools.html) has been updated to include a scope monitoring flag. Deployments that score Review Required on the five-criterion framework from WP2 but show no mechanism for tracking divergence between authorized scope and actual operating scope receive an elevated governance risk notation. The update reflects a consistent Stream Two finding: the intervention-point problem and the scope-drift problem are structurally connected. A system with well-designed intervention points but no scope monitoring will develop a governance gap as its operating scope expands beyond its original authorization.

Tool and working papers: <https://www.thinkcapital.org/research.html>

### FROM THE RESEARCH

#### Stream One:

In addition to the Working Paper 3 release referenced above, this past week's federal governance developments add a dimension the paper will address in a supplement: what governance frameworks calibrated to federal standards do during periods of federal policy uncertainty. Practitioners with state or local AI governance experience are invited to contact [michael.bragen@thinkcapital.org](mailto:michael.bragen@thinkcapital.org).

#### Stream Two:

The practitioner session in early June will focus on two questions emerging consistently from intake interviews: how agencies operationally distinguish between oversight of assistive AI and oversight of agentic AI, and how reviewer authority is defined within governance frameworks. Both questions connect directly to the NASCIO paradox described in this issue: an agency can have a responsible use policy that names a reviewer without having defined what the reviewer can do when they find a problem. The participation form is at <https://www.thinkcapital.org/research.html>.

## FIVE QUESTIONS FOR PRACTITIONERS

---

These questions apply across all three tiers examined this week. They identify where the gap between governance documentation and governance architecture is most likely to show up in practice.

1. Does your agency's AI governance framework function independently of any specific federal mandate? Is it calibrated to a reference that may shift? Can you identify what the framework governs if that reference is absent?
2. Does your current AI inventory reflect what your systems are doing today? Does it cover current integrations, data sources, and task delegation? Is the inventory updated periodically to cover extended scope beyond what systems were authorized to do at initial deployment?
3. For each named oversight role in your governance framework (reviewer, CPO, CAO, risk officer): does that person have the access, the defined criteria, and the operational authority to act on a finding? Or do they hold an advisory or documentation function?

4. Has your agency assessed whether its responsible use policy for citizen-facing AI resolves the specific concerns driving the 75% NASCIO figure? Does the policy address the documentation requirement in conjunction with addressing the operational risk?
5. How does your agency distinguish, in practice, between governance of assistive AI (analysis and recommendations) and governance of agentic AI (autonomous multi-step action)? Does the framework apply the same oversight model to both?

---

#### *Government AI in Practice*

*Published weekly by ThinkCapital LLC under the Government IT and AI Governance Initiative (GIAG). GIAG is a practitioner research program examining AI governance implementation in federal, state, and local government. Research participation, practitioner inquiries, and correspondence: [michael.bragen@thinkcapital.org](mailto:michael.bragen@thinkcapital.org). Archive and publications: [thinkcapital.org](https://thinkcapital.org).*

*WP3 (Mandate Translation): How Federal AI Governance Requirements Arrive at State and Local Agencies. Publication will be released in early June 2026 at [thinkcapital.org/publications.html](https://thinkcapital.org/publications.html)*

The views expressed are those of the researcher. Not for distribution without permission. Michael Bragen, Principal, ThinkCapital LLC | [michael.bragen@thinkcapital.org](mailto:michael.bragen@thinkcapital.org) | [thinkcapital.org](https://thinkcapital.org) | [thinkcapital.substack.com](https://thinkcapital.substack.com)

© 2026 ThinkCapital LLC. All rights reserved.