GOVERNMENT AI IN PRACTICE

Research and analysis from the ThinkCapital GIAG Initiative

ISSUE 1 · MARCH 2026

EARLY SIGNAL: FROM THE RESEARCH

Practitioners describing the most functional AI governance are not citing the RMF in day-to-day operations. They have translated framework language into agency-specific decision rules that their teams actually use.

The RMF is the architecture. What works operationally looks more like a field manual.

This distinction between framework as reference document and framework as operational practice is one of the central questions the GIAG research is designed to investigate comparatively across agency types.

From the Editor

Government AI adoption is accelerating. The research on whether it's working is not. This newsletter exists to close that gap.

Government AI in Practice publishes findings, practitioner perspectives, and evidence-based analysis from the ThinkCapital Government IT/AI Governance Initiative. Each issue combines original research with field observations from the people doing the actual work of implementing AI governance in federal, state, and local government settings.

If you've been inside one of these programs and are willing to share what you found, I want to hear from you. Details at thinkcapital.org/research.html.

Michael Bragen | Principal, ThinkCapital LLC

What We Don't Know About NIST AI RMF in Practice

The National Institute of Standards and Technology AI Risk Management Framework has become the closest thing the federal government has to a standard for responsible AI adoption. Agencies are citing it in acquisition documents, governance charters, and oversight committee briefings. That adoption is accelerating.

What is not accelerating is research on whether it works.

> ***Frameworks do not govern AI systems. Organizations do. And the gap between a well-designed framework and effective institutional practice is precisely where governance failures propagate.***

After spending 18 years conducting systematic assessments across 250+ organizations, including 30 government agencies, I have seen this gap before. The history of government IT is littered with frameworks that looked rigorous on paper and underperformed in practice, not because the frameworks were wrong, but because the organizational conditions required to implement them were never examined with the same rigor as the frameworks themselves.

Three Questions the Literature Cannot Answer

1. How does RMF implementation vary across agency types?

A defense intelligence unit and a state unemployment benefits agency both face AI governance requirements under the same federal guidance. Their risk profiles, workforce capabilities, mission constraints, procurement structures, and accountability architectures are fundamentally different. A governance approach that performs well in one context may be structurally incompatible with the other.

No published research has mapped that variation systematically. Agencies are making implementation decisions without comparative evidence on what conditions predict success in their specific context.

2. What separates durable governance from audit-cycle compliance?

Some organizations internalize governance frameworks as operational discipline. Others treat them as periodic documentation exercises. Both can produce paperwork that satisfies an oversight review. Only one produces governance that functions when it matters.

The conditions that differentiate these outcomes are empirically unexamined in the government AI context. Leadership investment, resource allocation, institutional culture, procurement structure, and workforce capability all plausibly matter. Which ones matter most, and under what conditions, is not known.

This distinction is not academic. When an AI system fails in a high-stakes operational context, the question of whether governance was operational or merely documented becomes very concrete, very quickly.

3. What does AI oversight actually look like when it is working?

Policy documents describe what human oversight of AI systems should look like. Practitioners describe something more complicated. Oversight mechanisms that function well in low-stakes environments can break down when decision timelines compress, workforce capacity is constrained, or AI outputs are difficult to evaluate without specialized knowledge.

The conditions under which human oversight remains meaningful in AI-augmented operations are not well documented. That gap matters most in exactly the agencies where the consequences of oversight failure are most severe.

Why This Research Matters Now

The window between widespread framework adoption and widespread framework evaluation is closing. NASCIO's 2026 state CIO survey found that between 80% and 90% of state CIOs have already implemented responsible use frameworks, and nearly 90% have created AI task forces. By every structural measure,

governance is in place.

And yet NASCIO's executive director identifies moving pilots to production as the major thrust for 2026. More than 1,000 state legislative AI bills in 2025 reflect genuine uncertainty about whether existing frameworks are adequate for what agencies are actually deploying.

Agencies that cite RMF compliance in acquisition documents today will face accountability questions about that compliance in oversight hearings tomorrow. CIOs making AI investment decisions right now are working from normative guidance rather than evidence on what has actually worked in comparable organizational contexts.

> ***The RMF is a reasonable starting point for government AI governance. What comes next requires knowing whether it is actually working.***

Michael Bragen is Principal at ThinkCapital LLC and founder of the Government IT/AI Governance Initiative. He spent 18 years at Software Productivity Research conducting IT assessments across 250+ organizations.

RESEARCH UPDATE | Government IT/AI Governance Initiative | March 2026

Where the Research Stands

The GIAG Initiative is conducting structured practitioner interviews with federal, state, and local government personnel who have direct experience implementing or evaluating AI governance frameworks. The research focuses on two primary streams:

Stream 1: NIST AI RMF Implementation

Examining how agencies across different contexts are interpreting and applying the RMF in practice. Key variables include agency type, risk profile, workforce capability, procurement structure, and the depth of governance beyond documentation compliance.

Stream 2: AI Oversight Quality in Government Operations

Examining what human oversight of AI-augmented decision-making actually looks like in operational settings, not what policy documents prescribe, but what practitioners describe when oversight is functioning as intended and when it is not.

The initiative is actively seeking practitioners with direct experience in federal, state, local, and tribal government contexts. Discussions can be on or off the record. Participants receive aggregate findings when the first research phase is complete.

To participate or learn more: thinkcapital.org/research.html

What to Watch

NASCIO Agentic AI Report (March 2026)

NASCIO's new report on agentic AI in state government raises governance questions the RMF was not designed to answer. When AI systems make sequential decisions faster than any human review cycle allows, what does meaningful oversight look like? Watch for OMB guidance on this in 2026.

State Legislative Activity

More than 1,000 state AI bills were introduced in 2025. The focus has shifted from regulating AI in government to regulating AI in the marketplace, with specific prohibitions on AI-driven eligibility determinations for Medicaid, healthcare, and insurance. The accountability implications for agencies already using AI in adjacent functions are not yet well understood.

The Pilot-to-Production Gap

NASCIO identifies moving pilots to production as the defining challenge for state CIOs in 2026. The governance implications are significant: frameworks designed for pilot environments may not be adequate for production systems operating at scale. This is an area the GIAG research is watching closely.

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