

---

# Government AI in Practice

Substack Posts Archive · Through April 2026

---

## Michael Bragen

Principal, ThinkCapital LLC

Government IT & AI Governance Initiative (GIAG)

[thinkcapital.substack.com](https://thinkcapital.substack.com) · [thinkcapital.org](https://thinkcapital.org)

[michael.bragen@thinkcapital.org](mailto: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).

---

**16**

Posts

**3**

Newsletter Issues

**2**

Research Streams

---

---

## Contents

[Apr 17, 2026](#) · *Newsletter Article*

Are You Governed or Just Documented?

[Apr 14, 2026](#) · *LinkedIn Commentary*

Procurement, Red Lines, and Operational Accountability

[Apr 14, 2026](#) · *LinkedIn Post*

Three Questions Before Your Next AI Deployment

[Apr 9, 2026](#) · *LinkedIn Post*

Scope Expansion as Incremental Accommodation

[Apr 8, 2026](#) · *LinkedIn Post*

When AI Becomes Load-Bearing Infrastructure

[Apr 1, 2026](#) · *Newsletter Article*

The M-25-21 Compliance Calendar Problem

[Apr 1, 2026](#) · *LinkedIn Post*

Classification Is Not Governance

[Mar 31, 2026](#) · *LinkedIn Post*

Governance Without Metrics Is Policy With Good Intentions

[Mar 31, 2026](#) · *LinkedIn Post*

Compliance Records vs. Compliant Systems

[Mar 27, 2026](#) · *LinkedIn Post*

You Cannot Manage What You Cannot Measure — AI Edition

[Mar 27, 2026](#) · *LinkedIn Post*

AI Governance Failures Are Operational, Not Policy

[Mar 25, 2026](#) · *LinkedIn Post*

Research Finding: Risk Registers Are Point-in-Time Artifacts

[Mar 23, 2026](#) · *Newsletter Preview*

The Pentagon-Anthropic Case Study: A Governance Architecture Lens

[Mar 19, 2026](#) · *LinkedIn Post*

Policy Velocity vs. Implementation Reality

[Mar 13, 2026](#) · *LinkedIn Post*

The RMF as Architecture, Field Manuals as Operations

---

---

April 17, 2026 · NEWSLETTER ARTICLE

## Are You Governed or Just Documented?

A research note from the GIAG Initiative.

Federal agencies can satisfy every NIST AI RMF documentation requirement and still operate a governance program that has never changed a single deployment decision.

That distinction is the central question driving the GIAG Stream One research. Two agencies can produce identical documentation portfolios and operate governance programs with entirely different functional characteristics. One program surfaces risk. The other records it.

The research is designed to investigate that gap — not at the policy level, where frameworks are reasonably well specified, but at the operational level, where implementation choices determine whether governance mechanisms function or simply exist.

[ Excerpt — full text available at [thinkcapital.substack.com](https://thinkcapital.substack.com) ]

---

April 14, 2026 · LINKEDIN COMMENTARY

## Procurement, Red Lines, and Operational Accountability

The procurement dispute between Anthropic and the Pentagon produced a useful clarification: AI developers can negotiate red lines into contracts, or they can rely on legal compliance as a proxy for safety constraints. Courts will eventually sort out which approach holds.

What neither approach addresses — and what this issue examines — is the layer beneath the contract: the operational environment where the model actually runs, where human oversight either functions or doesn't, and where accountability either has a clear owner or doesn't.

Contracts define what the vendor is responsible for. Governance frameworks define what the agency is responsible for. The space between those two things is where most deployment risk actually lives.

---

April 14, 2026 · LINKEDIN POST

## Three Questions Before Your Next AI Deployment

Issue 3 of *Government AI in Practice* publishes in two days. Here is one of the core diagnostic questions driving it.

Before your agency deploys its next AI system, answer these three questions:

1. Who updates the risk register when the model is retrained?
2. Who updates it when the deployment scope expands?
3. If neither question has a named owner with a defined trigger, the governance program is documented but not operational.

That is the distinction Issue 3 is built around.

---

April 9, 2026 · LINKEDIN POST

## Scope Expansion as Incremental Accommodation

Most AI scope expansion does not look like a decision. It looks like a series of reasonable accommodations.

A team finds a new use for a system that is already running. The expansion is small, the system handles it well, and no one flags it for review because no review trigger exists for incremental use case additions.

Six months later, the system is operating in a context materially different from the one that was assessed, authorized, and documented. The risk register still describes the original deployment.

This is not a compliance failure in the conventional sense. No one violated a policy. No one made a bad decision. The governance framework simply had no mechanism for capturing incremental scope expansion as a governance event.

---

April 8, 2026 · LINKEDIN POST

## When AI Becomes Load-Bearing Infrastructure

There is a specific moment in AI deployment when the governance calculus changes permanently. It happens when the system stops being a tool and becomes load-bearing infrastructure.

After that point, replacing or significantly modifying the system is more disruptive than managing the governance gap around it. The accountability conversation shifts from 'how do we govern this system' to 'how do we govern around this system.'

Most agencies do not have a formal process for identifying when that transition occurs. The system crosses the threshold incrementally, through accumulated dependencies, and the governance program does not register the change.

That transition point is one of the key phenomena the GIAG research is designed to identify and document comparatively across agency types.

---

April 1, 2026 · NEWSLETTER ARTICLE

## The M-25-21 Compliance Calendar Problem

The M-25-21 high-impact AI documentation deadline closed April 3. Agencies are reporting compliance. This issue examines what that compliance actually consists of, why the compliance calendar may be measuring the wrong thing, and what the asymmetric risk looks like for agency CIOs on both sides of the decision.

Also: a direct look at whether the designation process is generating governance change or documentation events — and why that distinction matters for every agency that has now classified a system as high-impact.

The deadline created a forcing function. Whether it created a governance program is a separate question, and the answer varies significantly by agency.

[ Excerpt — full text available at [thinkcapital.substack.com](https://thinkcapital.substack.com) ]

---

April 1, 2026 · LINKEDIN POST

## Classification Is Not Governance

Agencies racing to meet the M-25-21 deadline are treating the classification exercise as the governance action. Designating a system as 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 entirely.

An agency that has designated twelve systems as high-impact and updated its risk registers accordingly has done what the memo requires. Whether those risk registers are connected to operational decision processes — whether anyone reads them, acts on them, or updates them — is not something the memo measures.

---

March 31, 2026 · LINKEDIN POST

## Governance Without Metrics Is Policy With Good Intentions

The NIST AI Risk Management Framework has good intentions. The EU AI Act has good intentions. OMB M-25-21 has good intentions. None of them define what a functioning governance mechanism looks like in operational practice, or how you would know one if you saw it.

That is not a drafting failure. It is a measurement design gap. You cannot write a metric into a policy document that the implementing organization has not yet figured out how to collect.

The GIAG research is designed in part to fill that gap — to develop observational criteria for governance effectiveness that can be applied comparatively across agency implementations, without relying on self-reported compliance.

---

March 31, 2026 · LINKEDIN POST

## Compliance Records vs. Compliant Systems

A compliance record and a compliant system are not the same thing. Most federal AI governance audit mechanisms cannot tell them apart.

A system can produce a complete audit trail — documented risk assessments, signed authorizations, current risk registers — and still operate outside the parameters of its authorization. The documentation reflects the system as designed. The audit reflects the documentation.

That distinction is the central research question for GIAG Stream One: whether current federal AI governance frameworks have the structural capacity to detect operational drift, or whether they are designed to verify documentation rather than function.

---

March 27, 2026 · LINKEDIN POST

## You Cannot Manage What You Cannot Measure — AI Edition

Tom DeMarco wrote it in 1982: you cannot manage what you cannot measure. Capers Jones spent four decades building the empirical case. The functional sizing community gave practitioners the tools.

The principle has not changed. The technology has.

Government agencies are signing AI contracts right now without measurement frameworks to specify what they are buying, verify what was delivered, or detect when operational performance diverges from the contracted specification.

That is not primarily a technology problem. It is a measurement infrastructure problem — and it is exactly the kind of problem the functional sizing community has solved before, in different technology contexts.

---

March 27, 2026 · LINKEDIN POST

## AI Governance Failures Are Operational, Not Policy

Most AI governance failures are not policy failures. The policy language is there. The problem is operational.

An agentic system authorized as a document processing tool starts handling exception adjudication. A recommendation engine starts influencing resource allocation. Each expansion feels incremental. None individually triggers a formal review.

The governance framework was not designed to catch gradual scope expansion because no one anticipated that systems would expand that way. The authorization process is built around discrete deployment decisions, not continuous operational drift.

Fixing this requires rethinking the trigger architecture — what counts as a governance event, who is responsible for detecting it, and what the response mechanism looks like.

---

March 25, 2026 · LINKEDIN POST

## Research Finding: Risk Registers Are Point-in-Time Artifacts

Our latest research post is live on LinkedIn. Most agency AI risk registers accurately describe the system as designed at deployment. They do not describe the system as operating today.

That gap is not a documentation problem. It is a structural one. Agentic systems expand their operational footprint through use, through model updates, and through integration with adjacent systems — none of which automatically triggers a risk register update.

The result is a governance record that is technically complete and operationally stale. Auditors reviewing the record see compliance. Practitioners managing the system see something different.

---

March 23, 2026 · NEWSLETTER PREVIEW

## The Pentagon-Anthropic Case Study: A Governance Architecture Lens

The most important AI policy debate is not happening in Congress. It is happening in procurement.

Issue 2 of *Government AI in Practice* uses the Pentagon-Anthropic-OpenAI episode as the case study it deserves to be — not the politics, but the governance architecture.

When federal AI policy frameworks fail to specify what happens at the boundary between vendor safety constraints and agency operational requirements, that boundary gets resolved in contract negotiations, in court, or in deployment decisions made without adequate precedent.

The Pentagon procurement dispute is the first major public test of that boundary. Issue 2 examines what it reveals about the structural gaps in federal AI governance architecture.

---

March 19, 2026 · LINKEDIN POST

## Policy Velocity vs. Implementation Reality

OMB issued M-25-21 and M-25-22 in rapid succession in early 2025. Both memos carry specific AI governance requirements for federal agencies. Both arrived with implementation timelines that assumed agencies had baseline AI governance infrastructure already in place. Most did not.

The gap between policy issuance and implementation capacity is not new in federal IT. What is new is the velocity. AI policy is moving faster than agency governance infrastructure can absorb.

That velocity mismatch is producing a predictable outcome: compliance documentation that races ahead of governance capability. Agencies are meeting the deadlines. The question is what they are building behind the deadlines.

---

March 13, 2026 · LINKEDIN POST

## The RMF as Architecture, Field Manuals as Operations

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.

An agency that has built the field manual has done something the framework cannot do for it: translated governance intent into operational decision support. The research is designed to document what that translation looks like in practice, and what conditions make it more or less likely to occur.

---

## 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 April 27, 2026

---