GOVERNMENT AI IN PRACTICE

Research and analysis from the ThinkCapital GIAG Initiative

ISSUE 4 · LATE APRIL 2026

EARLY SIGNAL: FROM THE RESEARCH

Government agencies have built AI oversight programs designed for systems that recommend. They are now deploying systems that act. The governance architecture does not match the deployment.

The oversight problem is a planning failure. Technology deployment has outpaced governance architecture.

This issue documents the structural gap between what current governance frameworks require and what agentic AI deployments demand. It also examines what the EU AI Act requires that the NIST AI RMF does not, and what that gap means for government practitioners.

From the Editor

Issue 3 examined the governance timing problem: three clocks running behind the systems they are supposed to govern. The response from practitioners confirmed what the research was already finding. The timing gap is real, and agencies know it. What they are less certain about is the structural implication. Governance designed for advisory AI fails structurally when applied to agentic AI. The timing gap is real; the architecture gap is the deeper problem.

This issue moves from the timing problem to the architecture problem. When an agentic system executes an action, rather than producing a recommendation for human review, the oversight structure needs to be built differently. Most agencies have not done that work. Many have added AI officers, documented governance procedures, and satisfied compliance requirements. What most have not done is re-examine whether the oversight architecture they built for the previous generation of AI tools is adequate for systems that act autonomously in live workflows.

This issue also incorporates the EU AI Act and NIST AI RMF comparison that Stream Two research has been building. The comparison is not an academic exercise. US government practitioners are already encountering the gap, in international partnerships, in procurement negotiations with vendors operating across jurisdictions, and in advisory relationships with commercial clients who have EU obligations. Understanding what each framework requires, and where each is silent, is practical governance preparation. Practitioners encounter these gaps already.

A Wall Street Journal report published April 24 documents a pattern relevant to readers of this letter: multiple states have pursued AI transparency and safety requirements, and those efforts have stalled amid federal concerns about regulatory fragmentation. The structural tension is legitimate. Fragmented state rules create compliance burdens. The federal preemption argument has real force. It also has a critical dependency: a federal bill must actually pass. Until it does, AI systems continue to be procured and deployed across government contexts with no consistent accountability framework. Agencies absorb that risk whether they acknowledge it or not. The five intervention triggers GIAG Stream Two has identified, irreversibility, consequence transfer, distributional novelty, value conflict, and legal or regulatory significance, are designed to give practitioners a deployable accountability standard that functions in the current environment. No federal mandate is required. Issue 4 examines how agencies are building, or failing to build, the internal structures that make those triggers operational.

What Changed When AI Started Acting

The shift from advisory to agentic AI is primarily an operational change. In advisory deployments, the decision chain still runs through a human. A caseworker reviews the recommendation. A procurement officer approves the contract action. A compliance analyst confirms the flag before the file moves. Human review was built into the process by design.

Agentic systems alter that structure. They compress or eliminate the review interval. A system managing benefits eligibility routing does not pause to notify a supervisor before it redirects a case. A procurement assistant managing solicitation schedules does not wait for confirmation before advancing a milestone. The human reviewer is now downstream of the action, not upstream of it.

Three governance speeds matter here. Decision speed, the rate at which the system generates and executes actions, has accelerated substantially. Accountability speed, the rate at which responsibility for outcomes can be assigned, has not changed. Governance speed, the rate at which an agency can detect, assess, and respond to a problem in the system, remains the slowest of the three. Agentic AI widens all three gaps simultaneously.

The governance structures agencies built for AI recommendation systems are architecturally mismatched to AI action systems. That mismatch is where oversight fails.

Agencies designed oversight for the technology available when their governance architecture was built. The result is structures that fit the previous model and leave gaps in the current one. The question is whether those structures can be retrofitted or require a different starting point.

The Five Characteristics That Expose Nominal Oversight

GIAG Stream Two uses a decision-grounded framework to distinguish substantive oversight from documented oversight. Five characteristics of agentic AI decisions consistently signal where human review is necessary and where current governance falls short. The framework is developed in full in the Stream Two working paper, *When Humans Must Intervene*.

Characteristic	Government Oversight Implication
Irreversibility	Can the action be undone after execution? Most agencies have not mapped which agentic outputs are reversible. When a system routes a case, advances a deadline, or sends a notification, the reversal cost may be high or the window short. Oversight structures that assume reversibility as a default are inadequate for agentic deployments.
Consequence Transfer	Who bears the cost when the system errs? Government agentic AI regularly transfers consequence to the citizen: a delayed benefit, a misrouted appeal, a wrongly flagged compliance issue. Oversight designed around internal process metrics often misses this transfer entirely.
Distributional Novelty	Has the system encountered a situation outside its training distribution, or has operational drift carried it into novel territory? Most government agentic deployments lack systematic monitoring for distributional shift. The system keeps running; the novelty is invisible until something fails.
Value Conflict	When the system's objectives compete, who decides? Agentic systems operating at scale routinely surface situations where speed conflicts with accuracy, or consistency conflicts with equity. Without a designated human decision authority for these conflicts, they are resolved by the system's default weighting.
Legal / Regulatory Significance	Government action carries due process implications that agentic AI does not inherently respect. Decisions requiring a documented human judgment in manual processes require the same accountability standard when executed by an agentic system. Governance structures that treat automated outputs as process steps, not decisions, create legal exposure.

These five characteristics function as a diagnostic. Agencies that cannot answer basic questions about each one for their active agentic deployments have nominal oversight, regardless of what their governance documentation says.

Early Signal: What Stream Two Research Is Finding

Research Status: Stream Two practitioner interviews are in progress. Findings reported here reflect early-stage patterns from practitioner exchanges, structured conversations with government technology advisors, and the public deployment record. These are directional observations, not conclusions.

Oversight Was Designed for Procurement, Not Operations

The most consistent pattern across early Stream Two exchanges is that government agencies invested significantly in AI oversight during the acquisition and deployment approval phases, and substantially less in sustaining that oversight once systems went live. Governance structures tend to be concentrated at the front of the deployment lifecycle: the risk assessment, the authorization process, the initial policy documentation. Operational oversight, the ongoing monitoring of what agentic systems are actually doing, is thinner and often unassigned.

The problem is structural. The governance calendar excludes routine review of agentic system behavior after deployment. Staff responsible for initial compliance carry no responsibility for operational monitoring. The result is a credentialed system with degraded oversight.

AI Oversight Officers Exist. Escalation Authority Does Not.

Several agencies have designated AI officers or created governance committees in response to federal directives. What is less common is a defined escalation structure for agentic decision events: a named individual with the authority and the information to intervene when a system encounters a situation requiring human judgment.

Distributed ownership is the norm. Responsibility for agentic AI behavior is spread across the program office, the AI officer, the CISO, and the vendor. When something requires a human decision, the question of who decides is itself unresolved. The organizations handling this well are not fully centralizing governance. They are designating a named coordinator with defined override authority at specific escalation points.

'Human-in-the-Loop' Language Is Not Operationally Defined

Human-in-the-loop appears in nearly every agency AI governance document reviewed in Stream Two exchanges. The term is almost never operationally defined. It does not specify which decisions require human review. It does not specify who conducts the review, the review interval, the criteria for escalation, or the documentation standard. It functions as a policy commitment that satisfies a compliance requirement without creating an operational control.

The gap between what the governance document says and what the operational system does is where oversight fails. Stream Two research is examining this gap systematically.

EU AI Act vs. NIST AI RMF: What Each Framework Requires

Government practitioners working on agentic AI oversight are navigating two major frameworks that approach the problem from different directions and with different legal force. Understanding the gap between them is practical governance preparation.

The EU AI Act is a binding regulation with specific requirements for high-risk AI systems and, through its General Purpose AI model rules, for foundation-model-based deployments. The NIST AI RMF is a voluntary framework that establishes a risk management process without mandating what that process must achieve or what oversight must look like. For agentic AI, both frameworks leave significant gaps.

Dimension	EU AI Act	NIST AI RMF
Scope of obligation	Risk-category binary. High-risk systems carry mandatory requirements; everything else is largely unaddressed.	Risk-management process. All AI systems are in scope at some level; the rigor of application is calibrated to risk.
Human oversight requirement	Explicit and legally binding for high-risk systems. Technical measures must enable human oversight throughout the system's lifecycle.	Framework-defined but operationally undefined. Govern and Manage functions reference oversight; neither specifies what oversight must accomplish.
Agentic AI coverage	General Purpose AI Model rules (Article 51+) address foundation models but leave agentic deployment governance to member-state interpretation.	Silent on agentic-specific governance requirements. The RMF was not designed for systems that initiate actions autonomously.
Escalation / intervention standard	Requires that operators can intervene, override, and shut down. Specific technical capabilities required for high-risk categories.	No defined escalation standard. The framework does not specify what intervention capability must look like in practice.
Measurement and audit	Conformity assessment, post-market monitoring, incident reporting to national authority.	No mandatory measurement or audit. Agencies self-assess. No external verification standard for framework claims.
Practitioner implication	US agencies with EU-facing deployments or EU-based partners face binding requirements that exceed current federal guidance. Compliance planning must account for the gap.	Agencies operating under the RMF alone carry governance exposure that neither OMB guidance nor current audit practice is designed to surface.

Several implications are immediate for government practitioners.

First, US agencies with EU-facing deployments, international data-sharing partnerships, or vendor relationships that cross jurisdictions are already operating in environments where EU AI Act obligations apply, whether or not their domestic governance program accounts for them. This is a compliance gap that exists regardless of domestic policy positions. Procurement teams need to account for it now.

Second, the EU AI Act's human oversight requirements for high-risk systems are operationally specific in ways that OMB guidance is not. The system must be technically capable of human intervention, with functioning override mechanisms, named authority, and documented escalation paths. A policy statement asserting human-in-the-loop is insufficient.

Third, neither framework adequately addresses agentic AI governance as the deployment model is currently evolving. Agencies relying on either framework alone are building governance architecture against a prior-generation threat model. The current deployment environment requires practitioners to reason from first principles about where autonomous action ends and human accountability must begin, and to build that boundary explicitly into contracts, oversight structures, and operational policy before deployment, not after the first incident.

WHAT THE CURRENT FRAMEWORKS ARE MISSING ABOUT AGENTIC AI

Both the EU AI Act and the NIST AI RMF were developed with bounded AI systems in mind: systems with defined inputs, defined outputs, and a human review step. Agentic AI systems operate differently. They chain decisions, expand scope, and encounter situations their designers did not anticipate.

The EU AI Act's high-risk category rules require human oversight capability and intervention capacity, but **do not define what those requirements mean** when a system's scope drifts or when a multi-step agentic chain produces an outcome no individual step would have triggered.

The NIST AI RMF's Govern and Manage functions assume a governance team reviewing a bounded system. They **do not account for systems that continuously generate and execute actions** in workflows that evolve after deployment. Stream Two research is examining what governance architecture actually looks like in organizations that have confronted this directly.

Practitioner Diagnostic: Three Questions

The following three questions are drawn from the five-characteristic framework and from Stream Two practitioner exchanges. They are designed as a rapid self-assessment for government AI practitioners responsible for agentic deployments.

1	**Can you identify which agentic outputs in your deployment are irreversible?** If you cannot map reversibility, you cannot calibrate oversight. Irreversible actions require pre-execution review, not post-execution audit.
2	**Does your oversight structure include a named individual with defined override authority at designated escalation points?** If responsibility is distributed across multiple offices without a designated decision authority, you have diffusion, not governance.
3	**Is your 'human-in-the-loop' language operationally defined?** Specify: which decisions, who reviews, what interval, what documentation standard. If you cannot answer all four, the commitment is aspirational.

STREAM TWO WORKING PAPER

The five-characteristic framework referenced in this issue is developed in full in When Humans Must Intervene: A Decision-Grounded Framework for Human Oversight in Government and Commercial Agentic AI Deployments, available at thinkcapital.org/publications.html.

Our Participation Ask

Stream Two is actively recruiting government and public sector technology practitioners for structured research interviews. Participation involves a 45-minute conversation examining how your organization defines, assigns, and operationalizes human oversight in agentic AI deployments. Participation is confidential. Findings will be reported in aggregate and without attribution unless participants choose otherwise.

If you work in government IT, AI governance, or public sector technology advisory, and your organization has active or planned agentic AI deployments, this research is directly relevant to the problems you are working on. The conversation takes under an hour and is conducted under your choice of attribution terms.

Three practitioner types this research needs:

- Federal or state agency IT leaders, Chief AI Officers, or AI program managers with direct experience designing or implementing AI oversight programs.
- Government technology advisors or consultants who have advised agencies on NIST AI RMF implementation, EU AI Act compliance, or agentic AI governance architecture.
- Practitioners who have encountered the governance gaps described in this issue in their own organizations. The research needs the full range of experience, not only governance successes.

To express interest or schedule a conversation: thinkcapital.org/research.html | calendly.com/thinkcapital

What's Ahead

Issue 5 will report on the GIAG research intake: the participant profile taking shape across both streams, the implementation patterns emerging from the early interview set, and the first cross-agency observations on where governance variance is most pronounced. Structured interviews begin in May.

Working Paper 2, Human Oversight Quality in Agentic AI Deployments, has just been released and is available now at thinkcapital.org/publications.html. The paper presents the full five-characteristic framework for human oversight in agentic AI deployments with practitioner applications and sector-specific governance architecture recommendations for government and commercial contexts.

Government IT practitioners with experience governing agentic AI deployments are encouraged to reach out now. Stream Two interviews are open.

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, Belmont, California. 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