
ThinkCapital

Government IT/AI Governance Initiative

Research Participation Prospectus

A Comparative Study of AI Governance Framework Implementation and Oversight Quality in United States Government Operations

Version	1.0
Date	February 2026
Program Lead	Michael Bragen, Founder — ThinkCapital
Contact	researchprogram@thinkcapital.org
Web	thinkcapital.org/research.html
Status	Active — Participant Recruitment Open

This document is provided to individuals considering participation in the ThinkCapital Government IT/AI Governance research initiative. It describes the purpose of the research, what participation involves, how data will be handled, and what participants will receive in return. Please read it before confirming your interest in participating.

1. Purpose and Background

Government agencies across federal, state, and local jurisdictions are now operating under formal AI governance requirements. The NIST AI Risk Management Framework is accelerating in adoption. Oversight mechanisms are being documented, approved, and filed. Executive orders and OMB guidance have created a compliance infrastructure that did not exist three years ago.

What is largely absent from the literature is empirical research on whether any of it is working. Most governance commentary is normative: what agencies should do, what frameworks they should adopt, what policies should require. Research on what is actually happening in practice, and whether existing frameworks are producing the accountability outcomes they were designed to produce, is nearly nonexistent.

ThinkCapital is conducting structured research to address this gap directly. The initiative applies comparative assessment methodology developed through 250+ enterprise IT evaluations across federal agencies, defense organizations, and state governments to two specific and underexamined research questions.

This is independent practitioner research. ThinkCapital is not affiliated with any AI vendor, platform provider, or advocacy organization. It is not academic research requiring institutional review board oversight. The goal is empirical findings that working practitioners can use, not additional normative guidance for frameworks that already exist.

2. Research Streams

Stream One: NIST AI RMF Implementation in Government Practice

The NIST AI Risk Management Framework is the closest thing the federal government has to a standard for responsible AI adoption. Its design is thoughtful and its vocabulary is useful. But frameworks do not govern AI systems. Organizations do.

Primary research question:

"Where is NIST AI RMF adoption generating genuine risk reduction, and where is it producing compliance documentation without corresponding accountability?"

Key dimensions under investigation:

- How implementation varies across agency types, mission profiles, and accountability architectures
- What conditions predict durable governance versus audit-cycle compliance
- Whether documented RMF adoption reflects operational practice or compliance documentation
- How agencies adapt the framework to contexts it was not explicitly designed for
- The relationship between risk categorization accuracy and organizational incentive structures

Stream Two: Human Oversight Quality in AI-Augmented Government Operations

Every significant AI governance framework includes a human oversight requirement. What none of them define is what "meaningfully involved" actually requires in operational practice. That ambiguity creates a structural problem: agencies can satisfy oversight requirements by establishing review processes that are formally present but operationally inert.

Primary research question:

"What does meaningful human oversight actually require at scale, and how does current government practice compare against any reasonable operational definition of the term?"

Key dimensions under investigation:

- Oversight mechanism design across assistive and agentic AI system types
- Whether defined intervention points function as genuine controls or documented assumptions
- How oversight requirements are being applied to agentic systems that were designed for assistive AI
- Audit trail integrity and attribution mechanisms in AI-assisted decision workflows
- Conditions under which oversight requirements produce functional accountability versus rubber-stamp compliance

3. What Participation Involves

Participation is designed to be straightforward and respectful of your time. The core commitment is a single structured interview.

Format	One structured interview, conducted via video call or telephone
Duration	30 to 45 minutes
Scheduling	At your convenience; evenings and early mornings available
Questions	Provided in advance so you can prepare or decline specific topics
Recording	With your explicit permission only; transcripts shared back for review
Follow-up	Optional only; no obligation to respond to follow-up contact
Attribution	Your choice: named acknowledgment, organizational attribution, or full anonymity
Withdrawal	You may withdraw at any time, including after the interview, without consequence

We are particularly interested in practitioners with direct experience implementing or evaluating AI governance frameworks in the past 18 months. This includes those who encountered challenges, revised approaches mid-implementation, or found gaps between framework requirements and operational reality. Accounts of difficulty are as valuable to this research as accounts of success.

Participation in this research does not constitute an endorsement of ThinkCapital, its methodologies, or any specific governance framework or vendor.

4. Data Collection, Storage, and Security

Types of information collected:

- Professional role and title (for contextual analysis, not individual identification in published findings)
- Organization type (federal agency, state agency, local government, commercial, academic, industry association)
- Agency mission area or industry sector (broad categorization)
- Direct experience with AI governance framework implementation, evaluation, or oversight design
- Interview responses to structured questions — provided in advance for review
- Optional: organizational context that helps situate your experience in the comparative analysis

What we do not collect:

- Classified or sensitive operational information — questions are scoped to avoid this
- Information identifying specific individuals within your organization beyond your own role
- Financial, procurement, or contractual details
- Personally identifiable information beyond name and contact details provided voluntarily

Data storage and security:

Contact information and interview notes are stored in password-protected files accessible only to the ThinkCapital research team. Interview recordings, if made with your permission, are stored in encrypted cloud storage and deleted upon completion of transcript review. Raw interview data is retained for a period of three years following publication of findings, after which it is permanently deleted. You may request deletion of your data at any time by contacting researchprogram@thinkcapital.org.

Anonymization and aggregation:

Published findings will not identify individual participants unless explicit written permission is provided. Quotes used in publications will be reviewed by the participant before use. Findings are reported in aggregate, by organization type and sector, in a manner that does not permit identification of individual organizations or personnel.

5. What Participants Receive

Participation is not compensated financially. In exchange for your time, ThinkCapital commits to the following deliverables, provided at no charge:

Preliminary findings briefing	Access to preliminary findings before public release, with opportunity to provide review comments that may shape final analysis
Sector aggregate report	A summary of findings specific to your organization type (federal, state/local, commercial, academic) showing how your sector compares to the broader participant pool
Full published report	The complete final research report upon publication, including all methodology documentation and the comparative framework developed from findings
Participant acknowledgment	Named acknowledgment in published work, at your discretion — or full anonymity if preferred
Research updates	Periodic updates on the research program as it develops, including posting of intermediate findings and methodology notes

ThinkCapital does not sell, license, or share participant data with third parties, including AI vendors, government agencies, or commercial research organizations. Findings are published independently and are not commissioned by or provided exclusively to any sponsor.

6. Privacy Rights and Consent (GDPR and Equivalents)

ThinkCapital complies with the General Data Protection Regulation (GDPR) and equivalent privacy frameworks applicable to participants in their respective jurisdictions. The following rights apply to all participants:

Right of access: You may request a copy of all personal data ThinkCapital holds about you at any time.

Right to rectification: You may request correction of any inaccurate personal data.

Right to erasure: You may request deletion of your personal data at any time, including after the interview. Erasure requests will be honored within 30 days.

Right to restrict processing: You may request that your data be retained but not actively processed while a dispute or review is pending.

Right to withdraw consent: You may withdraw your consent to participate and to data processing at any time without consequence. Withdrawal does not affect the lawfulness of processing prior to withdrawal.

Right to object: You may object to specific uses of your data, including use in published findings, at any time prior to publication.

Right to data portability: You may request your data in a structured, commonly used format.

Consent statement:

By submitting the Research Engagement Form or confirming your participation in writing, you provide ThinkCapital with explicit consent to: contact you regarding research participation; collect and process the information you provide in the interview; retain that information for research purposes for a period of three years following publication; and use

anonymized or attributed content (per your stated preference) in published research findings.

This consent is specific to research participation and does not authorize use of your contact information for commercial, marketing, or any purpose unrelated to the Government IT/AI Governance Initiative. You may withdraw consent at any time by contacting researchprogram@thinkcapital.org.

7. Independence and Conflicts of Interest

ThinkCapital is an independent research and advisory organization. This research is self-funded and not commissioned by or conducted on behalf of any AI vendor, platform provider, government agency, or advocacy organization. No sponsor has editorial control over research findings, methodology, or publication decisions.

ThinkCapital does not receive referral fees, affiliate compensation, or any other consideration from vendors whose products or services may be discussed by research participants. The research program is designed to produce findings that are useful to practitioners regardless of which technologies or frameworks they have adopted.

8. Research Timeline

Q1 2026	Research Design and Recruitment Finalizing interview instruments, recruiting participants, establishing assessment criteria
Q2 2026	Primary Data Collection Conducting structured interviews, coding documentary evidence, iterative framework development
Q3 2026	Analysis and Draft Findings Cross-case analysis, pattern identification, draft framework development. Draft findings shared with participants
Q4 2026	Publication and Dissemination Final report published through ThinkCapital. Participant sector briefings. Submission to practitioners

9. Contact and Next Steps

To express interest in participating, complete the Research Engagement Form at thinkcapital.org/engage.html. The form asks for basic organizational context and your preferred research stream. You will receive an acknowledgment by email within one business day and a scheduling invitation within five business days.

For questions about the research program, methodology, or this prospectus, contact the research team directly:

Research email	researchprogram@thinkcapital.org
Program lead	Michael Bragen, Founder — ThinkCapital
Research program page	thinkcapital.org/research.html
Engagement form	thinkcapital.org/engage.html
LinkedIn	linkedin.com/in/michaelbragen

This prospectus is Version 1.0, dated February 2026. If you received this document through a third party, please verify currency at thinkcapital.org/research.html before submitting an expression of interest. ThinkCapital reserves the right to update this prospectus; material changes will be communicated to participants who have already expressed interest.

thinkcapital.org | researchprogram@thinkcapital.org | San Francisco Bay Area, California