What the GAAIA, the AI Action Plan, and CISA's Forthcoming Directive Mean Together

---

Abstract

Federal agencies now face simultaneous compliance demands from three instruments operating on independent timelines: *America's AI Action Plan* (July 2025), the draft *Great American Artificial Intelligence Act* (introduced June 2026), and a forthcoming CISA binding operational directive (BOD) on AI and cybersecurity. Each instrument is consequential individually. Together they create a governance requirement set whose interactions are not addressed by any of the three documents, and whose combined effect on agency planning is materially different from what any single instrument would produce alone. This article examines what the three instruments require, where they conflict or leave gaps, and what those gaps mean for federal and state agency technology leaders responsible for AI governance.

The analysis is directed at CIOs, CAIOs, CISOs, CDOs, and program managers in federal and state agencies that are deploying or procuring AI systems. It is grounded in GIAG research on NIST AI RMF implementation fidelity (Stream One) and human oversight quality in agentic AI deployments (Stream Two). The article identifies six specific open questions that remain unresolved across the three instruments, assesses the organizational risk each question creates for defined leadership roles, and provides a sequenced 120-day action framework for agencies working to close the gap between current governance posture and emerging compliance obligations.

Three Instruments, Three Legal Authorities, One Planning Problem

The *AI Action Plan* operates through executive authority. That distinction carries a specific operational implication: its directives are effective immediately and without legislative process, but they are also subject to revision or reversal by subsequent executive action. Governance programs built exclusively on executive directives carry inherent continuity risk. Agencies that anchor AI governance architecture to the Action Plan's institutional mechanisms (the CAISI role, the CAIO Council structure, or the

revised RMF content), are building on a foundation that a future administration could modify. That risk does not argue against compliance with current directives. It argues for building governance programs whose underlying logic is documented independently of the specific instrument that prompted them. Adaptation to future changes will require updating documentation rather than reconstructing the rationale from scratch.

The Action Plan sets direction on innovation, infrastructure, and international competition, and it establishes specific governance mechanisms: the Center for AI Standards and Innovation (CAISI) at NIST as the primary federal AI evaluation and compliance body, the Chief AI Officer Council as the interagency coordination venue, and a revised NIST AI RMF with specified content removed. These are current, active directives to federal agencies. CIOs and CAIOs who treat them as permanent structural features rather than current-administration posture are making planning assumptions that are risky from a historical perspective.

The GAAIA is proposed statute. It targets frontier AI developers with annual revenues above $500 million and significant compute operations, imposing mandatory risk assessments, third-party audits, and whistleblower protections. Civil penalties reach $1 million per violation per day. It designates CAISI as the compliance enforcement mechanism, allocating $300 million over three years. The proposed law preempts state AI laws for three years (plus a sunset provision). Currently the GAAIA is framed as a discussion draft. There is active opposition on the preemption question within both parties.

CISA's forthcoming BOD operates through existing regulatory authority and does not require Congressional action. It will create binding obligations for federal civilian agencies on AI security governance on its own timeline, regardless of the GAAIA's legislative trajectory.

The planning problem for agency CIOs and CAIOs is complex and urgent. The three instruments carry different legal weights, move on different timelines, and interact in ways that are not adequately addressed. Agencies are expected to build governance programs that satisfy all three simultaneously, against a set of open questions none of them resolves.

The RMF Revision Problem

The *AI Action Plan* directs NIST to revise the AI Risk Management Framework to remove references to misinformation, DEI, and climate change. Though that directive is unambiguous, what follows from it operationally is less clear. The revision timeline is unspecified. The content that replaces the removed material is unspecified. The transition guidance for agencies with existing governance programs built on RMF 1.0 documentation is unspecified. Federal and state agencies that used the current RMF as a compliance anchor now face a gap period: their governance documentation references an instrument under active revision, and they cannot update it to a framework that has not been published.

That gap carries audit risk. Inspector general reviews, FISMA assessments, and internal governance audits routinely reference RMF alignment as a compliance indicator. Agencies whose documentation cites specific RMF language that the revision removes will need either a transition strategy or an explanation of continued alignment under a changed framework.

The agencies best positioned to manage this are those that built governance programs around durable principles and documented their reasoning, rather than those that built compliance checklists tied to specific RMF document language. The former can adapt to a revised framework with relatively contained rework. The latter face the more difficult task of reconstructing the rationale behind governance decisions that were originally documented as RMF citations.

GIAG's Stream One empirical research is tracking this directly. The RMF revision creates a natural experiment in implementation fidelity: how agencies respond to a moving governance baseline reveals more about their actual governance maturity than their performance against a stable one.

CAISI as Enforcement Body: Capacity Against Assignment

The GAAIA assigns CAISI a compliance and audit function at significant scale, with a resource baseline of $300 million over three years. It is not a guarantee of operational readiness on any particular timeline, and the bill provides no milestones for when CAISI's enforcement infrastructure will be functional.

This creates a specific planning problem. Federal agencies are buyers of AI systems subject to GAAIA requirements. The bill's audit and risk assessment obligations imposed on developers will flow into federal procurement. Agencies will be evaluating vendor compliance claims referencing CAISI certification as part of acquisition governance. That evaluation requires knowing what CAISI certification means operationally (what standards were applied, what the audit scope covered, and what gaps a CAISI review does not address.) Those operational parameters do not yet exist. In short, agencies that build vendor accountability frameworks around a CAISI certification standard that has not been published are building on an undefined foundation.

The practical implication is that procurement governance work needs to proceed on two tracks: building the internal standards for what compliant vendor AI documentation looks like, and tracking CAISI's published guidance closely enough to calibrate those internal standards against the emerging external benchmark. Agencies that wait for CAISI to publish before beginning that internal work will find themselves behind when procurement decisions require it.

The CAIO Council Authority Question

The *AI Action Plan* formalizes the Chief AI Officer Council as the primary interagency AI coordination mechanism and directs integration with the CIO Council, CDO Council, Federal Privacy Council, and Chief Human Capital Officer Council. Although formalization is a meaningful structural step, the authority gap it leaves open is significant.

The Action Plan does not specify what decision authority the CAIO Council holds relative to CIOs and program managers in operational governance contexts. It establishes the Council as a coordination venue. Coordination without defined decision rights produces process overhead. When a CAIO and a CIO disagree about the governance requirements for a specific AI deployment, the Action Plan does not establish which function has authority to resolve that disagreement.

That ambiguity will matter in practice. AI governance decisions routinely implicate IT architecture, data management, procurement, security, and legal functions simultaneously. The CAIO's ability to drive consistent governance outcomes across those functions depends on a clear authority relationship with the officials who control them. Where that relationship is undefined, governance decisions default to whoever holds budget authority or program ownership, which is often the CIO or program manager rather than the CAIO.

It is imperative that agencies establish written internal authority protocols between the CAIO function and existing IT governance structures before external coordination requirements formalize around an unresolved gap. The CAIO Council formalization is an opportunity if agencies use it to clarify internal authority. It becomes a coordination burden if agencies wait for the Council itself to resolve the question.

The Preemption Uncertainty

The GAAIA's three-year preemption of state AI laws is the provision drawing the most public opposition. It is the provision most likely to change before any version of this bill is enacted. Active opposition comes from state lawmakers in Massachusetts and New York, AI safety advocates in both parties, and Republican governors including Ron DeSantis, who called Congressional AI preemption proposals "AI amnesty."

For agency technology leaders, the operational implication is straightforward. Agencies operating in California, New York, or Illinois cannot treat federal preemption as a planning assumption on the basis of a discussion draft. State AI laws in those jurisdictions are in effect now. The compliance obligations they create are current obligations.

The GAAIA's three-year sunset, if the bill is enacted with the preemption provision intact, functions as a rationalization window rather than a compliance holiday. It gives agencies time to build toward a unified federal compliance framework. Agencies must maintain state compliance programs in parallel throughout the legislative and implementation period.

Multi-state agencies face the most complex near-term picture. Their compliance architecture needs to accommodate both the possibility of federal preemption and the possibility that preemption fails or is significantly narrowed in conference. Building governance programs that are modular enough to adapt to either outcome is the more defensible planning posture.

The Agentic AI Oversight Gap

The AI Action Plan directs that all federal employees whose work could benefit from frontier LLM access should have it. The GAAIA's risk assessment requirements contemplate AI systems with significant operational autonomy. Neither document establishes a standard for what human oversight of agentic systems must look like to satisfy the compliance requirements both instruments create.

That standard will be built through agency implementation practice, CAISI guidance, and CISA enforcement precedent. The first agencies to define rigorous agentic oversight architectures will influence what becomes the de facto federal standard, ahead of formal regulatory prescription.

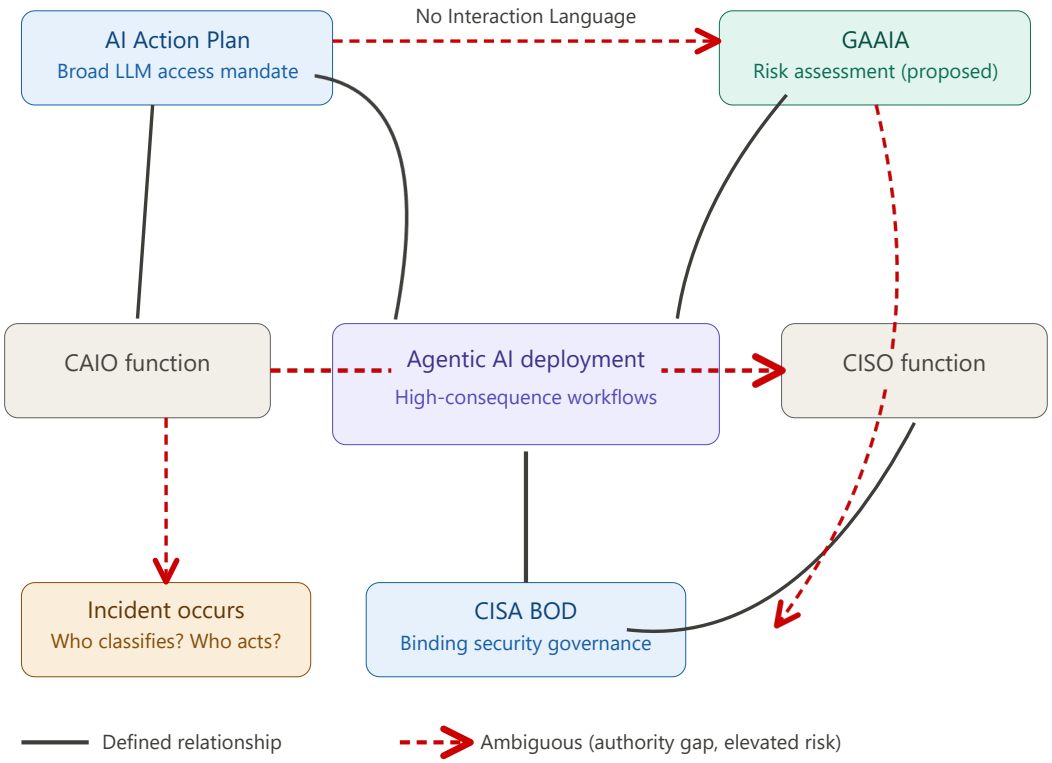

The CISA playbook directive is the most concrete near-term signal. The Action Plan directs modification of CISA's incident response playbooks to require CISOs to consult with CAIOs on AI-related incidents. That requirement acknowledges a boundary problem that is directly relevant to agentic deployments: in an autonomous AI workflow, the distinction between a cybersecurity incident and an AI governance failure is often ambiguous at the moment the incident occurs. A model behaving unexpectedly in a high-consequence process could be a security event, a governance failure, a data quality problem, or some combination. Classifying it correctly requires expertise and coordination that most agencies have not yet institutionalized.

The CISO-CAIO consultation requirement creates the coordination mandate. It does not resolve the decision authority question when those two officials reach different conclusions about incident classification or response. Agencies that document explicit protocols for that scenario before the BOD publishes will be ahead of every agency that builds them reactively.

GIAG's Stream Two research is examining human oversight quality in agentic AI deployments directly. The research question the current document set sharpens is this: as the federal government accelerates agentic AI adoption under the Action Plan's broad access mandate, what oversight mechanisms are

being designed, and what is their actual effectiveness at the consequential decision boundaries where autonomous action and human review must interface?

Open Questions for Agency Leadership

The following six questions are operationally significant for agency planners. None of them is resolved by the three documents under review. Each creates a specific exposure for one or more leadership roles.

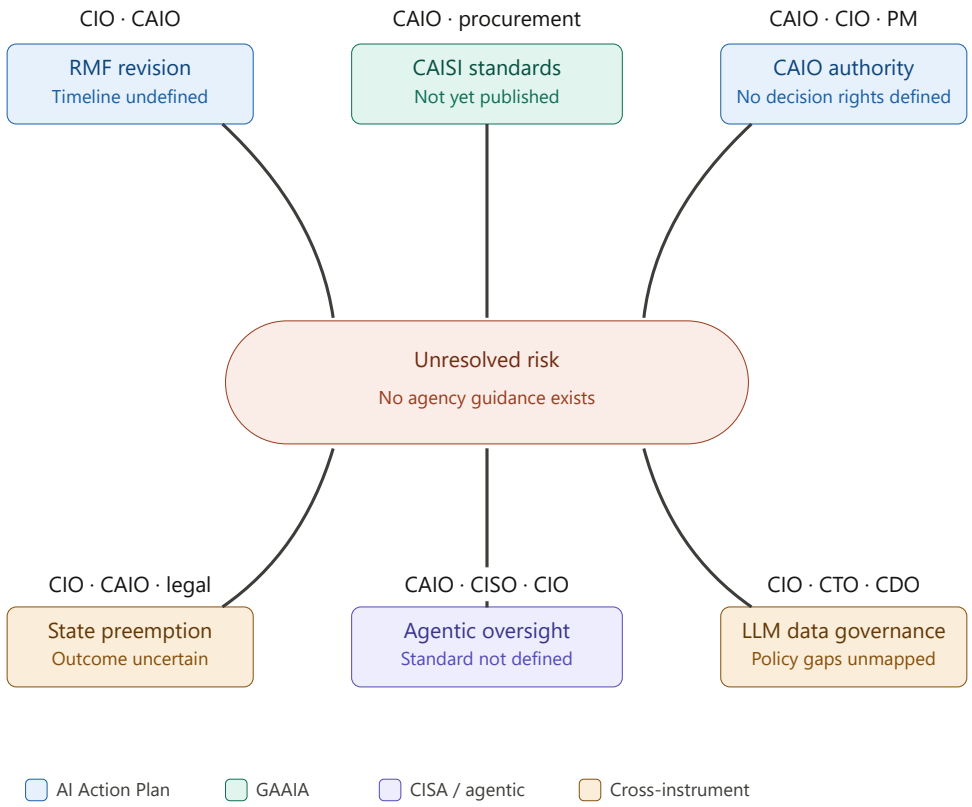

Open questions across the three instruments. See text for a full analysis

RMF revision timeline and transition guidance. NIST has not published a revision timeline, and the Action Plan provides none. For CIOs and CAIOs, this creates an immediate documentation integrity problem. IG reviews and FISMA assessments reference RMF alignment as a compliance indicator. If an agency's AI governance program cites specific RMF 1.0 language that the forthcoming revision removes, the documentation will be misaligned with the current governing instrument before any remediation action is possible.

Agencies need to know whether NIST will provide a formal transition period, what replaces the removed content, and whether existing program documentation will be grandfathered or must be updated on a compliance timeline. Until those answers are available, CIOs should treat any RMF-anchored documentation as provisional and flag it accordingly in governance records.

CAISI vendor compliance certification standards. The GAAIA designates CAISI as the compliance enforcement body for frontier AI developers. In practice, this means that federal procurement officers and CAIOs will need to evaluate vendor documentation that claims CAISI certification. That evaluation requires knowing what CAISI certification actually covers: what standards were applied, what the audit scope included, what was excluded, and what remediation was required. None of those parameters exist yet. The practical risk is specific: an agency awards a contract to a vendor presenting CAISI-certified risk assessment documentation, accepts that certification at face value, and later discovers the certification covered model development practices but not deployment-phase oversight or data handling in the agency's specific operational context. The contracting officer and CAIO bear accountability for an assumption neither had the tools to test. Procurement governance work on internal AI vendor evaluation standards should begin now, not after CAISI publishes.

CAIO Council decision authority relative to CIOs and program managers. The Action Plan formalizes the CAIO Council as the primary interagency AI coordination mechanism but does not define what the CAIO can direct versus recommend at the agency level. In practice, AI governance decisions routinely cut across IT architecture (CIO), data management (CDO), security (CISO), acquisition (procurement officer), and legal functions simultaneously. When a CAIO determines that a deployed agentic system requires additional oversight controls and the program manager or CIO disagrees on scope or timeline, the unresolved authority question defaults to whoever controls the budget. That is rarely the CAIO. CAIOs without documented decision authority will find themselves issuing governance recommendations that program managers can accept or defer without consequence. Agencies should establish written internal authority protocols between the CAIO and CIO functions before the CAIO Council formalization creates external coordination requirements that expose the internal gap.

GAAIA preemption outcome and multi-state compliance architecture. The three-year preemption provision is the most contested element of the GAAIA and the one most likely to be modified or removed before enactment. State AI laws in California, New York, and Illinois impose obligations on AI developers and, in some provisions, on deploying organizations. Federal agencies with operations in those states, grant-making functions that touch state and local recipients, or shared-services arrangements with state agencies are not operating in a preemption-protected environment on the basis of a discussion draft. CIOs and CAIOs in those contexts need a compliance architecture that accommodates both preemption passing and preemption failing. The modular planning posture is to document state compliance obligations separately from federal compliance obligations, maintain both in parallel, and design governance programs so that harmonization is achievable rather than requiring reconstruction if the legal landscape shifts.

The agentic oversight standard under concurrent GAAIA and CISA requirements. The GAAIA requires risk assessments for AI systems above the covered threshold. CISA's BOD will impose security governance requirements on agentic deployments. Neither document defines what constitutes adequate human oversight of an agentic system operating in a high-consequence government workflow. That definition matters because it determines accountability. If an agentic procurement system makes a vendor selection recommendation that a contracting officer approves without meaningful review, and that decision later produces an adverse outcome, the question of whether adequate human oversight occurred will be evaluated against whatever standard CAISI and CISA establish. CIOs and CAIOs who are deploying agentic systems now are operating ahead of that standard. The agencies that document their oversight design decisions and the reasoning behind them will be better positioned when the

standard is published than those that cannot reconstruct how (or why) their oversight architecture was designed.

Workforce-scale LLM deployment and data governance compatibility. The Action Plan's directive to provide frontier LLM access to all employees whose work could benefit is the broadest deployment mandate in the document. It creates immediate pressure on data governance frameworks that were designed for controlled, role-specific AI access. The specific risks are not abstract. An employee using a frontier LLM to process procurement-sensitive data may inadvertently expose information outside the intended access boundary. A benefits adjudicator using an LLM assistant to summarize case files may introduce model error into a decision with legal consequences for the claimant. Neither scenario requires malicious intent. Both require data governance and oversight frameworks that define what data can enter the LLM environment, what outputs require human verification before use, and what audit trail the interaction must leave. CIOs who have not mapped their existing data classification and access control architecture against a workforce-scale LLM deployment scenario are carrying unquantified exposure. CTOs who are selecting or building the underlying infrastructure need those data governance parameters before deployment architecture decisions are finalized, not after.

Where Things Stand and What Needs to Happen

Some agencies are ahead of this curve. Those with mature IT program management offices, documented data governance frameworks, and established CAIO functions with clear internal authority have the organizational infrastructure to absorb simultaneous governance requirement changes without significant disruption. They are updating documentation, calibrating procurement standards, and building CISO-CAIO coordination protocols as deliberate program activity. They represent a meaningful minority of the federal and state agency population.

Most agencies are not in that position. AI systems are in production. Governance documentation is anchored to an RMF under revision. CAIO functions exist on paper with authority that has never been tested against a real governance dispute. Vendor AI compliance is being accepted on self-attestation. Data governance frameworks designed for structured, role-specific access are being applied to workforce-scale LLM deployments they were not designed to cover. This is not a criticism of individual leadership. It is a structural observation: the pace of AI deployment has outrun the pace at which governance infrastructure can be built inside organizations that were not designed for this kind of simultaneous multi-front compliance pressure.

The cumulative risk is not theoretical. Consider the scenario an IG or GAO reviewer would construct twelve months from now: an agency deployed an agentic AI system for benefits adjudication under the Action Plan's broad access mandate. The system was procured from a vendor whose compliance documentation was accepted at face value because CAISI standards had not yet published. The agency's AI governance program was documented against RMF 1.0 language that the revision subsequently removed. When an adverse outcome occurred, the CAIO issued a corrective recommendation that the program manager deferred, because no written authority protocol established the CAIO's decision rights. The state in which the agency operates had an active AI law that federal preemption, which failed in conference, was assumed to have resolved. None of these failures required bad intent. Each one was

the predictable consequence of building AI governance programs on undefined foundations, and then deploying production systems before the foundations were resolved. That scenario is not a worst case. It is a plausible near-term outcome for agencies that do not begin closing these gaps now.

The practical challenge is that none of this work arrives with dedicated resources or a clear organizational home. Documentation reviews compete with operational demands. CISO-CAIO coordination protocols require two executives with crowded calendars to agree on language neither has a template for. Procurement standard development requires legal, acquisition, and technology functions to collaborate on requirements that current acquisition vehicles were not written to accommodate. Authority charters for the CAIO function require senior leadership engagement at a moment when most agency heads are focused on budget constraints and political priorities. The work is necessary and unglamorous. Agencies that treat it as a governance program management activity, assigning ownership and tracking completion, will make measurable progress. Those that treat it as background noise will find themselves in reactive compliance posture as enforcement mechanisms mature.

Priority actions for the next 60 days. The most urgent work is documentation integrity and incident protocol. CIOs should conduct a targeted audit of existing AI governance documentation to identify every citation to NIST AI RMF 1.0 language. The immediate deliverable is a registry of those citations, flagged as subject to revision, with the underlying governance rationale documented independently of the specific RMF language. This is not a full rewrite. It is a scoping exercise that establishes the rework backlog before NIST publishes the revision and the backlog becomes a compliance timeline. Agencies that complete this audit are in a position to update documentation quickly when the revised RMF publishes. Those that have not audited their citations will be estimating rework scope under time pressure. Simultaneously, CAIOs should draft a written CISO-CAIO coordination protocol covering AI incident classification, escalation criteria, and decision authority when the two officials reach different conclusions.

The document should fit on two pages. It needs to exist before CISA publishes the BOD. An absent protocol at the time of a CISA review is a finding. A protocol that exists and was followed is evidence of governance maturity, even if it requires subsequent refinement. CTOs deploying or evaluating agentic AI infrastructure should map current access control and data classification policies against the workforce-scale LLM deployment scenario and document where the policies are silent or insufficient.

Priority actions for the 60-to-90-day window. The second priority is procurement governance and agentic oversight readiness. CIOs and CAIOs should jointly define the minimum acceptable content for AI vendor compliance documentation, independent of what CAISI eventually publishes. This standard should specify the required scope of a vendor risk assessment: system boundary definition, deployment context coverage, audit methodology description, disclosed limitations, and explicit statements about what the assessment did not evaluate. This document becomes the agency's internal evaluation baseline for current and near-term AI procurements. Contracts awarded without this standard in place are being evaluated against an undefined criterion. When CAISI publishes its standards, this internal document becomes the calibration point for alignment, not the starting point for thinking about the problem. For agencies with agentic AI systems in production or in procurement, this window is the right time to conduct an oversight architecture gap analysis: mapping current human review points against the consequence-tier of each AI-assisted decision, documenting where human oversight is nominal rather than substantive, and identifying the highest-risk gaps for remediation. That gap analysis serves two

purposes: it is a governance risk management tool now, and it is a preliminary compliance readiness assessment if the GAAIA advances. Agencies that have documented their oversight architecture and its gaps are in a fundamentally different position than those that cannot describe their oversight design when a regulator asks.

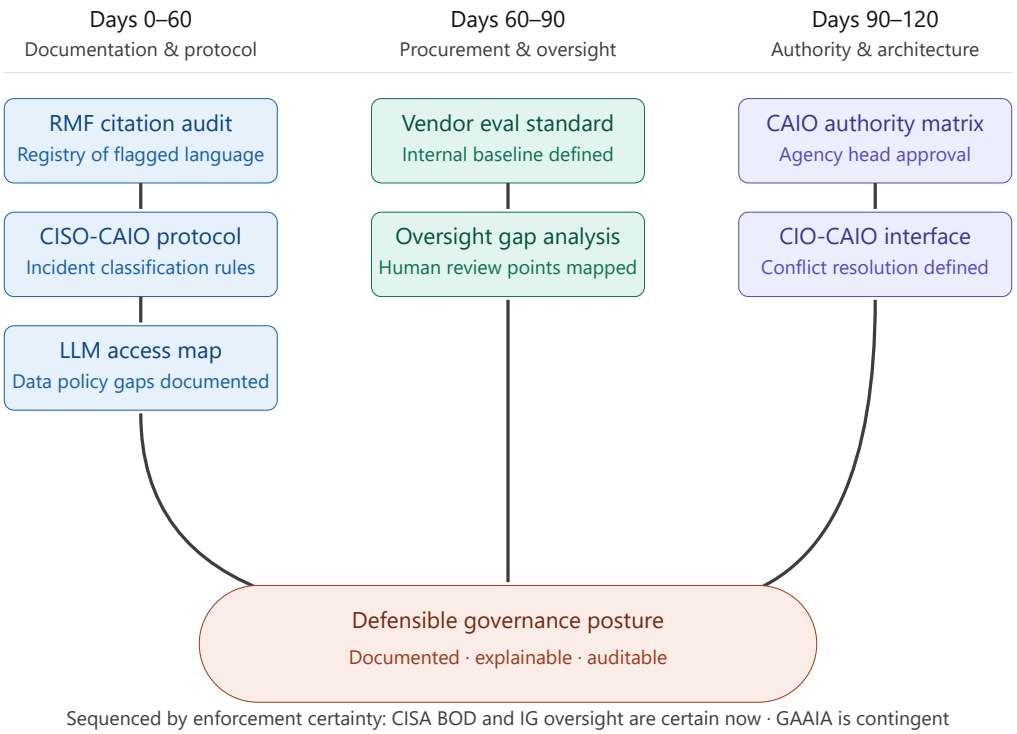

Priority actions for the 90-to-120-day window. The third priority is authority structure and governance architecture. CAIOs should use the CAIO Council formalization as the occasion to establish a written authority matrix at the agency level: which AI governance decisions the CAIO can make unilaterally, which require CIO or CDO concurrence, which require CISO review, and which require escalation to the agency head or deputy. Without this document, the CAIO function remains advisory in the disputes that matter most. The authority matrix should be approved by the agency head or deputy, not just drafted by the CAIO, because its value depends on organizational legitimacy, not just the CAIO's preference. CIOs should treat the interface between their IT governance frameworks and the CAIO function as an architectural gap that requires a documented solution, not an informal working relationship that functions until it does not.

The specific question to resolve is what happens when a CIO-controlled program decision and a CAIO governance requirement conflict. Leaving that question to be resolved under pressure during an incident is the organizational equivalent of the absent incident protocol: a gap that becomes consequential at the worst possible moment. Agencies that have completed the documentation audit, the procurement standards definition, the oversight gap analysis, and the authority matrix by the end of this window have a defensible governance posture against the three instruments reviewed here. The posture is not perfect. The instruments are still evolving and several open questions remain unresolved. But the agency will be able to demonstrate to an IG reviewer, a CISA auditor, or a program oversight committee that

governance decisions were made deliberately, documented carefully, and built on a foundation the agency can explain and defend.

GIAG's Stream One research is generating practitioner data on RMF implementation fidelity as the governance baseline shifts. Stream Two is examining how human oversight mechanisms are being designed in agentic deployments ahead of formal standard-setting. Both streams are producing findings directly relevant to the planning and risk questions this article identifies.

ThinkCapital LLC provides advisory services to federal and state agencies across the full range of governance readiness work described above: documentation audit and transition planning, AI vendor compliance standard development, agentic oversight architecture design, CISO-CAIO coordination frameworks, and CAIO authority structure. Structured assessment frameworks and advisory engagements are available for agencies at any stage of this work, from initial gap assessment through implementation support. Practitioners working through these issues or navigating them in federal or state agency contexts are welcome to engage directly.

Direct research and study engagement inquiries to: research@thinkcapital.org