DESIGN THE OPERATING MODEL

The Operating Model for AI at Scale

By Brittney Hannah and Cedric Kwindja

The question federal agencies now face is not whether to deploy AI. Most already have — in pockets, in pilots, in isolated programs that produced promising results and then stalled. The question is whether AI will remain a series of experiments or become something the organization actually runs on.

That distinction is structural. It cannot be resolved by technology alone.

What separates agencies that achieve sustained AI performance from those that cycle through vendor engagements without durable change is not the sophistication of the models they deploy. It is how they govern, organize, and execute AI as an enterprise function. It is the operating model.

The evidence is accumulating. Documented AI use cases across major federal agencies nearly doubled from 571 in 2023 to 1,110 in 2024. Generative AI usage alone increased ninefold over the same period.[1] And yet, as of 2024, the majority of those use cases remained in a preliminary phase — initiated, but not operationalized.[2] The capability is present. The operating infrastructure to sustain it is not.

Why the Operating Model Is the Real Work

Federal leaders often arrive at AI strategy having already made significant technology investments. What they discover — sometimes late — is that the technology was never the binding constraint.

The binding constraints are organizational: unclear ownership, fragmented decision rights, workforce roles that were never redefined, and governance structures borrowed from IT project management that were never designed for AI's pace of change or its cross-mission reach.

An AI operating model addresses these constraints directly. It defines who owns AI outcomes across the enterprise, how decisions get made and by whom, what the relationship is between central AI functions and mission-facing components, and how the organization learns and adapts as AI capabilities evolve.

Without this structure, even technically excellent AI deployments degrade. Models drift. Use cases stall waiting for approvals no one is clearly authorized to give. Mission components build redundant capabilities in parallel. And the organization never accumulates the institutional knowledge that compounds over time into genuine strategic advantage.

The DoD's experience with its Chief Digital and Artificial Intelligence Office illustrates the dynamic precisely. A delayed chartering directive and unresolved boundaries with the DoD Chief Information Office created persistent confusion about roles and authorities — not because the capability was absent, but because the governance architecture was incomplete.[3] A function without clear decision rights cannot

govern. It can only advise. And advisory functions do not scale AI.

The Four Structural Questions

Designing an AI operating model for a federal agency requires answering four foundational questions. The answers will look different across agencies — they should. But every organization operating AI at scale must have considered each one seriously.

Who owns AI outcomes?

Ownership is not the same as accountability on a RACI chart. It means someone is positioned, resourced, and authorized to ensure AI delivers mission value — not just that AI systems are running. In practice, this requires an enterprise AI function with genuine authority: not advisory, not advisory-adjacent, but empowered to set standards, allocate resources, and enforce coherence across components.

Many agencies have created Chief AI Officer roles without creating the organizational conditions for those roles to function. A title without decision rights is a constraint, not a capability. Where those roles carry real authority — backed by clear policy, budget influence, and executive access — AI programs move. Where they do not, programs stall at the boundary of the next component that did not sign up for the enterprise vision.

How are decisions made — and at what level?

AI introduces a new category of decisions that traditional federal governance structures were not designed to handle: decisions that are cross-functional, move faster than the budget cycle, and have downstream consequences that are difficult to anticipate. The operating model must define which decisions belong at the enterprise level, which belong at the component or program level, and which can be delegated further.

This is not about centralization for its own sake. Centralized governance of the wrong decisions creates friction without adding control. The goal is to centralize what must be consistent — standards, risk thresholds, data access governance, model evaluation criteria — and decentralize what must be responsive to mission context.

DHS's approach to the authorization problem is instructive here. Rather than forcing AI deployments through traditional ATO cycles that were never designed for iterative AI development, DHS shifted toward a continuous authorization model — a secure-by-design approach that gives AI developers access to existing secure systems without waiting on static approval timelines.[4] That is a governance decision, not a technology decision. It resolved a structural friction point that was throttling deployment speed.

What roles does the organization actually need?

Federal AI deployments frequently underinvest in the roles that sustain AI performance: AI product owners who understand both the mission and the model, data stewards who maintain the integrity of the inputs AI depends on, and decision-support specialists who help mission personnel use AI outputs appropriately. These are not IT roles. They are mission roles with an AI dimension, and most workforce plans have not

caught up to that reality.

Agencies that cite insufficient technical talent as a barrier to AI maturation are, in many cases, not describing a hiring problem. They are describing a role-design problem.[5] DHS recognized this directly, committing to a tailored staffing strategy to identify technical and multidisciplinary skill needs and embed them across components — not concentrate them in a central office where they cannot reach the mission.[6]

The operating model must specify not just the technical roles — data scientists, ML engineers, AI architects — but the operational roles that connect AI capability to mission execution.

How does the organization learn?

AI systems are not static. Models degrade, mission requirements shift, and new capabilities emerge faster than most federal planning cycles accommodate. An operating model that does not include a formal mechanism for monitoring performance, incorporating feedback, and evolving governance as conditions change will be obsolete within eighteen months of deployment.

This is one of the places where federal agencies face the most structural friction. The mechanisms for sustained learning — continuous evaluation, rapid iteration, adaptive governance — run against the grain of how most federal programs are funded and managed. Addressing it requires deliberate design, not optimism.

Three Operating Model Configurations

There is no single correct structure for federal AI governance. The right configuration depends on mission complexity, the current state of AI maturity across components, and the degree of interoperability required across the enterprise. Three configurations are most common in mature federal AI programs — each with a real federal analog.

Federated with enterprise standards. A central AI function sets policy, standards, and risk governance. Individual components retain implementation authority within those standards. DHS operates closest to this model. Its components — CBP, TSA, FEMA, Secret Service — have missions too distinct to govern from a single center, so the enterprise function sets the guardrails and the components execute within them. The risk, visible in DHS's own governance experience, is that standards become compliance exercises rather than genuine operational guidance. Sustaining the model requires the enterprise function to remain engaged at the component level, not just at the policy level.

Centralized with mission support teams. Core AI capabilities are built and maintained centrally, with embedded support teams that work alongside mission components. The DoD's CDAO was built on this logic — establishing enterprise scaffolding including federated model catalogs, MLOps infrastructure, and shared data services that components could draw on rather than rebuild independently.[7] This concentrates expertise and accelerates capability development, but requires the central function to remain genuinely responsive to mission needs. Where the CDAO encountered friction was in maintaining that responsiveness while managing the scope of an enterprise-wide mandate. The structural lesson: the embedded support function must be as strong as the central platform.

Hybrid with center-of-excellence architecture. A center of excellence operates as an internal consulting function — providing specialized expertise, setting architectural standards, managing enterprise-level programs — while components develop their own capabilities within a defined framework. HHS's OneHHS model reflects elements of this approach: an AI-integrated Commons that connects sub-agencies for data sharing and joint development, designed to allow new AI solutions to be developed, tested, and deployed across different operational environments without each component starting from zero.[8] With over 270 documented AI systems spanning NIH, FDA, CMS, and dozens of other offices,[9] HHS could not centralize execution — but it could create the connective tissue that makes enterprise-level coherence possible.

The selection of a configuration is a strategic decision, not a technical one. It has implications for budget authority, workforce planning, and the pace at which AI capabilities can be scaled. It should be made deliberately, with clear criteria for how the model will evolve as the organization matures.

What Sustainable AI Execution Requires

An operating model is a structure. Execution is what converts structure into performance. Federal agencies that operate AI sustainably share several characteristics that go beyond organizational design.

They treat AI governance as a mission function, not a compliance function. The distinction matters. Compliance-oriented AI governance asks: are we within policy? Mission-oriented AI governance asks: is AI producing outcomes that advance what we are here to do? HHS articulated this directly in its AI strategy, framing governance not as a control layer but as the mechanism by which the department builds public trust in AI-enabled service delivery.[10] The first question is necessary. The second is what drives performance.

They invest in the human infrastructure that AI depends on. Data quality, workforce capability, and change management are not prerequisites that get resolved before AI begins. They are ongoing operational requirements that the operating model must account for continuously. The agencies reporting the greatest depth of AI integration — HHS, VA, DHS — share a common pattern: they have institutionalized these investments rather than treating them as one-time program costs.

They build feedback loops that are fast enough to matter. Federal program cycles are measured in years. AI performance cycles are measured in weeks. The operating model must create mechanisms — formal and informal — that surface performance signals quickly enough for decision-makers to act on them. This often requires changes to how programs are structured and funded, not just how AI is governed.

They hold the line on accountability. AI creates diffuse accountability by design — multiple systems, multiple vendors, multiple components all contributing to an outcome. The operating model must counteract this by assigning clear, named accountability for AI outcomes at every level. Not shared accountability. Named accountability.

The Transition from Capability to Competitive Position

Federal agencies that get the operating model right do not simply manage AI better. They begin to operate differently — with greater speed, precision, and strategic reach than agencies that have not made the structural investment.

This is the distinction that matters most for mission leadership to understand. AI as a capability is something that can be purchased, deployed, and, if it fails, replaced. AI as an operating advantage is embedded in how the organization makes decisions, allocates resources, and pursues mission outcomes. It compounds over time. It creates institutional knowledge that does not leave when contracts end.

The agencies furthest along — not just in use case count, but in operational integration — are those that resolved the authority question first. They determined who owns AI outcomes and gave that function the organizational standing to enforce standards, before trying to scale capability. The ones still cycling through pilots without enterprise traction have governance structures that look complete on paper but lack real decision rights in practice.

The transition from capability to advantage is not primarily a technology question. It is a leadership question, an organizational design question, and a governance question.

It is exactly the work the operating model is designed to do.

The Sainth works with federal agencies to design the organizational structures, decision rights, and governance frameworks that allow AI to operate at scale — sustainably, and with strategic intent. Our work is not advisory in the conventional sense. We are embedded partners in the design of how agencies operate and compete.

For leaders ready to move from AI experimentation to enterprise-wide advantage, we are prepared to begin.

Notes

1. U.S. Government Accountability Office, *Artificial Intelligence: Generative AI Use and Management at Federal Agencies*, GAO-25-107653, July 29, 2025.

2. Nextgov/FCW, "Agency AI use doubled in 2024, GAO finds," July 29, 2025.

3. Federal News Network, "CDAO provides AI implementation plan, policy directive to DoD IG," December 2024.

4. Department of Homeland Security, *DHS AI Strategy for OMB Memorandum M-25-21*, September 2025.

5. Nextgov/FCW, "Agency AI use doubled in 2024, GAO finds," July 29, 2025.

6. Department of Homeland Security, *DHS AI Strategy for OMB Memorandum M-25-21*, September 2025.

7. Chief Digital and Artificial Intelligence Office, *Department of Defense Compliance Plan for OMB M-24-10*, U.S. Department of Defense.

8. Nextgov/FCW, "HHS releases AI strategy, united by new OneHHS approach," December 5, 2025.

9. Nextgov/FCW, "Agency AI use doubled in 2024, GAO finds," July 29, 2025.

10. Nextgov/FCW, "HHS releases AI strategy, united by new OneHHS approach," December 5, 2025.