FOUNDATIONAL PAPER NO. 2 · TECHNOLOGY

AI and the Federal Workforce

Artificial intelligence is the first technology in generations that could make government simultaneously cheaper and better, or more expensive and less accountable. The workforce question decides which.

CENTER FOR AI & PUBLIC SECTOR INNOVATION

JULY 2026

Three thousand six hundred and eleven. That is the number of individual artificial intelligence use cases federal agencies reported to the Office of Management and Budget for 2025, more than double the prior year's total.[1] Fifty-six agencies submitted inventories. More than 1,800 of the reported systems are already deployed or in pilot. The Department of Health and Human Services alone reported 447 use cases; the Department of Veterans Affairs reported 367, including more than 200 designated high-impact because they touch decisions about individual citizens.[2] Whatever the federal government's posture toward artificial intelligence was supposed to be, its actual posture is now a matter of public record. The technology is inside the building.

The question that matters has therefore changed. It is no longer whether government will adopt AI but whether the two-million-person workforce operating these systems is prepared to command them. That question will be answered either deliberately or by drift, and the taxpayer funds both answers. This paper surveys the current state of federal adoption, identifies the workforce transformation it demands, assesses institutional readiness honestly, and closes with recommendations for an AI-ready workforce that improves service while shrinking cost.

I. The State of Adoption: From Pilot Theater to Policy

Federal AI policy has moved through three phases in five years. The first, opened by Executive Order 13960 in December 2020 and codified by Congress in the fiscal 2023 defense authorization, established the annual public inventory of agency AI use cases.[3] The second phase layered on risk governance. The third and current phase, set by Executive Order 14179 in January 2025 and implemented through OMB Memorandum M-25-21, reorients the government toward accelerated adoption: agencies must designate Chief AI Officers, expand use of lower-risk AI, apply minimum risk-management practices to high-impact systems before deployment, and conduct honest maturity

assessments.[4] A companion memorandum, M-25-22, overhauls how agencies buy AI, with explicit attention to avoiding vendor lock-in.[5]

The inventories reveal a government adopting faster than it is transforming. The doubling from 1,757 reported use cases in 2024 to 3,611 in 2025 reflects genuine deployment growth, and it also reflects the pattern the numbers cannot hide: adoption is concentrated. Five large agencies have accounted for over half of all reported use cases for three consecutive years, while small agencies collectively contributed two percent of the 2025 inventory.[6] The distribution matters because the citizen does not experience the federal government at its median. He experiences it at a specific counter, in a specific queue, at whichever agency holds his claim, and the gap between the government's most and least capable agencies is now a gap in the quality of citizenship itself.

Two agency examples carry the pattern. At Veterans Affairs, AI is concentrated where the mission meets the individual veteran: claims and clinical workflows where a triage model can compress a backlog that human review alone could never clear, which is why VA leads the government in high-impact designations.[2] At Health and Human Services, the largest inventory in government spans fraud detection, biomedical research, and benefits administration.[2] In both cases, the honest description of most deployed systems is augmentation: the AI reads, sorts, drafts, and flags, and an accountable human decides. That division of labor is the correct one, and preserving it as systems grow more capable is a workforce problem before it is a technical one.

II. The Workforce Question

Every technology sorts work into three piles: tasks it automates, tasks it augments, and tasks it newly creates. For the federal workforce, the automation pile is dominated by high-volume document processing, data entry, first-line correspondence, and routine eligibility screening, functions that employ large numbers of the government's clerical and administrative staff. The augmentation pile is larger and more consequential: claims examiners, contract specialists, auditors, caseworkers, and analysts whose judgment remains the product but whose reading, drafting, and searching can be radically accelerated. The creation pile is the one government is least prepared for: model evaluators, AI procurement specialists, data stewards, and the supervisory tier that must learn to audit a machine's output the way a good editor audits a draft.

The arithmetic of the augmentation pile is where the taxpayer's interest concentrates. If AI tools return even a modest fraction of the working week to a claims examiner or a contract specialist, the compounded effect across two million employees is measured in tens of thousands of work-years annually. Captured, that surplus becomes shorter queues and smaller headcount through attrition. Squandered, it becomes what productivity gains become in institutions without performance discipline, which is invisible slack. The difference between capture and squander is management, measurement, and training, three words that describe a workforce agenda rather than a technology agenda.

> *A decision the citizen cannot appeal to a responsible official is a decision improperly made.*

One principle must govern the entire transition. Accountability remains with humans. M-25-21 embeds this in its minimum practices for high-impact AI, and the reason is older than the technology: in a republic, the exercise of public power must always trace to an official who can answer for it.[4] A decision the citizen cannot appeal to a responsible official is a decision improperly made, whatever its accuracy. The workforce implication is direct. The government must train not merely operators but adjudicators, people who understand a system's failure modes well enough to overrule it, and who retain the standing to do so.

III. Readiness, Honestly Assessed

Institutional readiness has three gaps, and candor requires naming each.

The skills gap. The government has never centrally tracked its training investment with the rigor the private sector applies, and what evidence exists suggests systematic underinvestment: training budgets are historically the first casualty of appropriations pressure, and the government's own strategic human capital management has sat on the Government Accountability Office's High-Risk List since 2001, with critical skills gaps cited as a core reason.[7] An institution deploying 3,611 AI systems on a workforce it has chronically undertrained is running an experiment on itself.

The procurement gap. The Clinger-Cohen Act gave agencies modern IT acquisition authority in 1996, and three decades of experience since have demonstrated that authority without expertise produces expensive failure.[8] AI procurement raises the stakes: agencies must evaluate model claims they cannot independently benchmark, negotiate data rights they may not recognize as valuable, and avoid the vendor lock-in that M-25-22 explicitly warns against.[5] The scarcest resource in federal AI is a contracting officer who understands what he is buying.

The governance gap. Chief AI Officers and governance boards now exist on paper across the agencies. Whether they mature into genuine institutional capacity or calcify into compliance ritual depends on staffing them with people who hold both technical fluency and operational authority, a combination the government's classification and pay systems were never designed to recruit.

IV. The Strongest Objection

The most serious objection to an acceleration agenda deserves full statement. It holds that government is precisely the wrong institution to move fast with this technology. Federal decisions are coercive and consequential in ways private decisions are not; a hallucinated paragraph in a marketing draft costs a company embarrassment, while an erroneous automated judgment in a benefits or enforcement context costs a citizen his due process. Government's poor record with large

technology programs suggests it will buy badly, deploy carelessly, and then defend its failures with the self-preserving instinct of every bureaucracy. Prudence, on this view, counsels waiting until the technology stabilizes and others have absorbed the early errors.

The objection is right about the stakes and right about the record, and its conclusion does not follow. Waiting is not a neutral act, because the status quo is itself a running harm: the backlog that ages a veteran's claim, the fraud that flows through unexamined payment streams, the fourteen-month permit. Each is a cost of the government's present processing capacity, paid continuously by identifiable citizens. Nor does delay purchase safety. Employees at agencies without sanctioned tools are already using unsanctioned ones, and shadow adoption is adoption without governance, the worst available combination. The choice on offer has never been between risk and no risk. It is between governed deployment under M-25-21's accountability practices and ungoverned drift. The conservative disposition, rightly understood, prefers the arrangement in which a responsible official can be identified, and that is the governed one.

V. Recommendations: An AI-Ready Workforce

First, train at the scale of the deployment. Every employee whose work touches a deployed AI system receives role-specific training in its capabilities and failure modes, with agency training investment reported annually as a per-employee figure the public can compare. What the government spends per employee on reskilling should be a published number, because what is unmeasured will be cut.

Second, build the adjudicator tier. Formal certification for employees who supervise high-impact AI, covering model limitations, escalation, and override authority, so that the human in command is a trained fact rather than an organizational chart fiction.

Third, professionalize AI acquisition. A government-wide cadre of AI-literate contracting professionals, shared across small agencies that cannot sustain their own, with data rights and lock-in review as standard checklist items under M-25-22.

Fourth, convert productivity into published savings. Agencies deploying AI at scale should report where the recovered work-years went: backlog reduction, cycle-time improvement, or headcount restraint through attrition. The taxpayer who funds the transition is entitled to see its yield.

Fifth, keep accountability human and visible. Every high-impact system publicly inventoried, every consequential decision traceable to a responsible official, every citizen's avenue of appeal preserved as a matter of design.

The efficiency tradition in American government runs from the two Hoover Commissions through the Grace Commission to Clinger-Cohen, and its lesson is consistent: technology multiplies the institution it enters.[9] A disciplined workforce becomes more capable; an undisciplined one becomes expensively undisciplined. Artificial intelligence will not decide which kind of institution the federal

government is. The people who train, buy, govern, and command it will. Companion studies in this institute's Backgrounder series take up particular applications of these tools; this paper's claim is prior to all of them. Fund the transition deliberately, because the taxpayer will fund it either way.

NOTES

1. Office of Management and Budget, 2025 Federal Agency Artificial Intelligence Use Case Inventory (released Apr. 2026), consolidating 3,611 reported use cases across 56 agency submissions, up from 1,757 individually reported use cases in 2024.
2. OMB, 2025 Federal Agency AI Use Case Inventory: agency totals for the Department of Health and Human Services (447) and the Department of Veterans Affairs (367, including more than 200 high-impact designations); more than 1,800 systems reported deployed or piloted government-wide.
3. Exec. Order No. 13960, Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government (Dec. 3, 2020); James M. Inhofe National Defense Authorization Act for Fiscal Year 2023, Pub. L. No. 117-263, § 7225 (codifying agency AI use case inventories).
4. Exec. Order No. 14179, Removing Barriers to American Leadership in Artificial Intelligence (Jan. 23, 2025); OMB Memorandum M-25-21, Accelerating Federal Use of AI Through Innovation, Governance, and Public Trust (Apr. 3, 2025).
5. OMB Memorandum M-25-22, Driving Efficient Acquisition of Artificial Intelligence in Government (Apr. 3, 2025).
6. Brookings Institution analysis of the 2024 and 2025 federal AI use case inventories (Apr. 2026), reporting concentration of use cases among the five largest reporting agencies.
7. U.S. Government Accountability Office, High-Risk Series: strategic human capital management, designated high risk since 2001, with mission-critical skills gaps as a continuing basis for the designation.
8. Clinger-Cohen Act of 1996, Pub. L. No. 104-106, div. E, 110 Stat. 679 (codified in relevant part at 40 U.S.C. § 11101 et seq.).
9. Commission on Organization of the Executive Branch of the Government (First Hoover Commission), Reports (1949); Second Hoover Commission, Reports (1955); President's Private Sector Survey on Cost Control (Grace Commission), Report (1984).