# Contextual AI for Federal Enterprise Digital Enablement

How contextual AI can accelerate U.S. Government modernization, service delivery, and operations while improving public outcomes, trust, and workforce effectiveness

## Enterprise Lifecycle Innovation and Excellence


**Contact Information**
Elevate Federal, LLC
6185 Walters Way
Foresthill, CA 95631
UEI: DSXLM7FLLLP3
CAGE: 11WG0
Primary NAICS: 541611
Socio-economic: WOSB
**info@elevatefederal.com**
**(888) 748-5550**

**Prepared by**
David Skoglund, COO
david.skoglund@elevatefederal.com
(571) 655-1874



**Publication Date**
April 2, 2026


WOSB CERTIFIED

ELEVATE Federal

# Contents

## Executive Summary

Federal enterprise digital enablement initiatives are trying to move agencies from fragmented, project-by-project modernization toward reusable, measurable, and outcome-oriented delivery. Across the U.S. Government, that means faster paths from mission need to trusted release, better public and workforce experiences, stronger control over cost and risk, and more consistent use of reusable platforms, shared services, and delivery patterns. The same need appears in benefits administration, health and human services, contact centers, grants, inspections, taxation, identity, and mission-support shared services.

Contextual AI is the strongest way to accelerate those objectives. Unlike generic AI, contextual AI works against authoritative agency context: current policy and forms, workflow states and exception paths, enterprise architecture patterns, data definitions, prior decisions, operational signals, accessibility requirements, and real user research. Used this way, AI becomes more than a writing assistant or code generator. It becomes a structured accelerator for intake, requirements, fit-gap analysis, decomposition, testing, documentation, support knowledge, and operational learning.

| **Working positions for leaders and evaluators** |
|---|
| 1. Enterprise digital enablement should change the delivery operating model, not merely add more disconnected tools or projects. |
| 2. Generic AI is not enough for government. Contextual AI—grounded in authoritative policy, data, workflow, architecture, user research, and control context—is the higher-value approach. |
| 3. AI should compress translation work, rework, and manual coordination while keeping humans responsible for judgment, approval, exceptions, and public-facing accountability. |
| 4. The best early targets are intake, requirements, fit-gap analysis, decomposition, configuration, testing, knowledge support, and operational learning—not unsupervised autonomy in consequential decisions. |

The right model is human-accountable, not human-absent. AI should reduce translation work, rework, and staffing intensity in repeatable delivery lanes, while humans retain decision rights, exception handling, and last-mile finishing. For high-impact uses that affect care, benefits, rights, safety, or public trust, AI should assist and structure judgment—not silently replace it. [1][2][3][4][5]

| What enterprise digital enablement is trying to accomplish | What good looks like in practice |
|---|---|
| Speed with control | Less time from request to trusted release, without sacrificing evidence, review, or auditability. |
| Reusable delivery | More buying, configuring, integrating, and reusing; less reinvention of one-off patterns. |

| | |
|---|---|
| Mission and customer alignment | Work tied to constituent, beneficiary, business, employee, and public journeys rather than technology for its own sake. |
| Lower delivery friction | Fewer handoffs, better backlog quality, clearer ownership, and less swivel-chair effort across teams. |
| Better economics | Lower unit cost, less duplicate tooling, and more transparent tradeoffs among build, buy, and modernization options. |
| Trust and compliance | Privacy, security, civil rights, records, accessibility, and explainability designed in from the start. |
| Scalable adoption | Patterns that can be applied across administrations, programs, field locations, and—in modified form—other agencies. |

Source basis for public-policy framing: [1] OMB M-25-21; [2] OMB M-25-22; [3] NIST AI RMF and Generative AI Profile; [4] Executive Order 14058 and Federal CX guidance; [5] 21st Century IDEA / GSA digital experience guidance.

## 1. Why federal enterprise digital enablement initiatives emerge

Digital enablement programs usually emerge when the existing delivery model is no longer keeping pace with mission needs. Across the federal public record, the pressures are visible: long-running legacy modernization challenges, mission-critical IT programs that run late or underperform, uneven requirements and stakeholder engagement, fragmented workflows, technical debt, digital-experience compliance gaps, acquisition friction, and the need to scale AI under stronger governance and public-trust expectations. Federal policy now adds urgency by directing agencies to accelerate AI adoption around innovation, governance, accountability, and mission outcomes. [1][4][5][6][7][8][9][10]

| Public signal | What public sources suggest | Implication for enterprise delivery and AI |
|---|---|---|
| Repeated modernization difficulty | GAO continues to describe long-running federal legacy-modernization challenges, high operations-and-maintenance burdens, and schedule risk in critical systems. | Do more shaping and risk reduction earlier, before scale. AI should help structure discovery, options, and evidence, not just write code. |
| Requirements and stakeholder gaps | GAO's high-risk and mission-critical IT work continues to point to weaknesses in upfront planning, acquisition discipline, stakeholder engagement, and development practices. | AI should be applied to intake, elicitation, synthesis, decomposition, and acceptance criteria, with human validation. |
| Workflow fragmentation | Customer-experience and digital-experience guidance emphasize end-to-end journeys, seamless service delivery, | AI must understand real workflow states, exception paths, and cross-system dependencies, not only idealized process maps. |

| | | |
|---|---|---|
| | and fewer channel breaks across federal services. | |
| Cost, procurement, and platform friction | Government-wide modernization still competes with high operations-and-maintenance spending, acquisition lead times, and economically constrained build/buy/configure choices. | AI-assisted decisions about build, buy, configure, migrate, or retire, should include economic, acquisition, and platform context. |
| Governance and trust expectations | OMB and NIST now require agencies to pair innovation with governance, inventories, proportionate risk practices, and trustworthy AI controls. | Speed must be paired with documented controls, role clarity, and evidence that AI use is safe, reviewable, and accountable. |

Challenge sources: [1] OMB M-25-21; [6] GAO-25-107795; [7] GAO-25-106908; [8] GAO-25-107852; [9] GAO-24-106764; [10] GAO-26-107859.

## Public success inputs that support this direction

- Technology Modernization Fund case studies show the value of reusable platforms, shared services, service-centered modernization, and standardization that can be adapted across agencies. [11]
- Login.gov shows the power of reusable shared services that improve access, security, consistency, and privacy protections across agencies and jurisdictions. [12]
- IRS Direct File and other tightly scoped digital-service efforts show that user-centered modernization can improve clarity, completion, and operational learning when teams iterate quickly against real feedback. [12]
- Government-wide CX, digital experience, acquisition, and AI policy materials show that the federal environment is building the management, literacy, and oversight foundations needed to scale trustworthy AI more broadly. [1][2][4][5][13]

## 2. Contextual AI: the missing design requirement

Most AI acceleration efforts feed the design model a user story, some source documents, and a code repository. That is nowhere near enough to meet the vision of government and private industry leaders alike. To reduce time to mission use, and still improve quality, maintainability, and trust, teams need a context fabric around the model. In government, the failure mode of generic AI is usually not lack of fluency; it is lack of authorized context.

Contextual AI means the model works inside approved environments and is grounded in authoritative sources, current operating rules, and real service conditions. For AI to be successful, it should understand not only what the department or agency wants to build, but also how it is allowed to work, what systems already exist, where exceptions occur, what constituents,

businesses, grantees, or federal staff actually experience, and what evidence will be needed for release, support, and oversight.

| Traditionally underused context inputs | What is often missed | Why it materially improves AI output | Examples of outputs improved |
|---|---|---|---|
| Policy, statute, regulation, and adjudication rules | Teams provide summaries but not the authoritative source language, decision rules, or current revisions. | Reduces policy drift, strengthens requirements quality, and limits confident but noncompliant output. | Requirements, decision logic, controls, notices |
| Forms, letters, templates, and plain-language standards | AI sees narrative goals but not the exact artifacts citizens and staff use. | Improves fit to real operations, reduces rewrite cycles, and yields clearer public-facing content. | Forms, letters, web content, scripts |
| Workflow states, routing rules, and local exceptions | Idealized process maps omit escalation paths, work queues, and site-level realities. | Prevents brittle automation and better reflects how work is really completed. | Process maps, next-step suggestions, SOPs |
| Enterprise architecture and approved platform patterns | AI is asked to design or code without knowing target platforms, integration patterns, or guardrails. | Produces more maintainable and scalable solutions, with fewer dead-end designs. | Solution options, integration patterns, scaffolding |
| Authoritative data models and business definitions | Terminology is inconsistent across programs, systems, and reports. | Improves semantic accuracy and reduces rework caused by mismatched definitions. | Data mappings, APIs, prompts, analytics |
| Security, privacy, records, and access controls | Teams add controls late, after requirements or prototypes are already drafted. | Avoids rework, reduces approval friction, and improves traceability for risk decisions. | Control matrices, design notes, workflows |
| Accessibility, Section 508, and user research | Accessibility and user research are treated as downstream checks instead of core inputs. | Leads to more usable, inclusive, and plain-language outputs from the start. | UI text, flows, acceptance criteria, test cases |
| Legacy system behavior, defects, and change history | AI sees target-state aspirations but not recurring root causes and brittle dependencies. | Improves modernization realism and helps avoid repeating old failure patterns. | Migration plans, refactor options, test seeds |
| Operational telemetry, tickets, call summaries, and training gaps | Designs ignore what service desks, call centers, and field users already know. | Turns lived operational pain into structured requirements and better support content. | Backlogs, KB articles, feedback summaries |
| Cost, licensing, and acquisition context | The model optimizes technical elegance without understanding economic or contractual constraints. | Supports better build/buy/configure choices and more realistic roadmaps. | Option analysis, sequencing, retirement plans |

Policy and governance context for this section: [1][2][3][4][5]. This table extends those public principles into a practical operating model for delivery teams.

## 3. High-value AI application lanes

The best early use cases are not the most autonomous ones. They are the places where delivery teams spend time translating one artifact into another: mission need into backlog, policy into decision logic, workflow reality into fit-gap analysis, architecture standards into implementation patterns, and production signals into the next improvement cycle. These are the lanes where contextual AI can reduce manual effort while still leaving humans in charge of judgment and release.

| Application lane | What AI should do | Human last mile | Primary outcome |
|---|---|---|---|
| Enterprise intake and prioritization | Classify requests, identify reuse, surface dependencies, draft an initial play decision and scope frame. | Mission owners confirm priority, risk, and budget posture. | Faster intake and better portfolio discipline |
| AI-assisted Sprint Zero and requirements | Generate interview prompts, summarize workshops, derive epics and stories, draft acceptance criteria, and expose assumptions. | Product owners and SMEs validate scope, wording, and edge cases. | Higher-quality backlog and fewer downstream surprises |
| Process discovery and fit-gap | Compare current workflow to desired-state platform options, identify bottlenecks, and highlight exception paths. | Business and operations leads decide which gaps matter and which processes should change. | More realistic modernization decisions |
| Configuration, low-code, and integration scaffolding | Draft configuration patterns, interface notes, mappings, field logic, and first-pass technical scaffolding. | Architects and engineers approve patterns, security fit, and maintainability. | Faster build with stronger standardization |
| Test, accessibility, and quality evidence | Generate test seeds, trace requirements to tests and controls, draft accessibility checks, and assemble evidence packages. | QA, accessibility, security, and release leads verify evidence independently. | Better release readiness with less manual compilation |
| Knowledge management and casework support | Turn policy and procedure sources into grounded retrieval experiences, staff guidance, summaries, and response drafts. | Frontline staff retain decision authority and tailor the final response. | Faster service and lower search burden |
| ATO, compliance, and control support | Crosswalk design decisions to controls, privacy artifacts, records requirements, and approval packages. | Risk, privacy, and security officials review the evidence and accept or reject risk. | Lower approval friction and clearer audit trails |

| | | | |
|---|---|---|---|
| Operations and continuous learning | Summarize telemetry, tickets, call trends, user feedback, and defects into actionable backlog recommendations. | Leaders decide what to fix, retire, or scale next. | Faster learning loops and better reuse decisions |

These lanes align with the federal push for AI that increases quality of public services, government efficiency, and reuse, while keeping risk management proportionate to impact. [1][2][3]

## 4. The operating model changes

If contextual AI is used well, the staffing model shifts. The change is not that expertise disappears; it is that routine synthesis, cross-walking, documentation drafting, first-pass scaffolding, and repetitive evidence compilation consume far fewer hours. The remaining work becomes more specialized and more consequential: product management, service design, architecture, domain stewardship, integration judgment, privacy, security, accessibility, records, testing strategy, and final approval.

In practical terms, this means smaller cross-functional squads for bounded delivery lanes, fewer handoffs between analysis and build, faster decision support for modernization choices, and stronger emphasis on human last-mile finishing. The winning pattern is neither human-in-every-loop nor human-in-no-loop. It is human-in-the-right-loop.

### AI-adapted sprint cycle

1. Intake and framing: AI assembles the initial context pack from policy, forms, prior work, architecture standards, tickets, and dependencies.
2. Elicitation and decomposition: AI drafts questions, summarizes workshops, generates epics and stories, and proposes acceptance criteria and trace links.
3. Optioning and build: AI supports fit-gap analysis, drafts configuration or code scaffolds, and assembles test and control artifacts.
4. Independent verification: humans review content, verify system state and evidence, and approve only what meets policy, quality, and mission thresholds.
5. Release and learn: AI turns telemetry, feedback, and support signals into the next prioritized improvement set.

| Risk tier | Typical use | Human role | Default posture |
|---|---|---|---|
| Assist | Search, summarization, retrieval, draft artifacts, and knowledge support. | Human reviews before use in an official process or public-facing action. | Broadly acceptable starter lane |
| Recommend | Triage, prioritization, fit-gap suggestions, next-best-action proposals. | Human approves the recommendation and owns the decision. | Good for medium-risk productivity gains |

| | | | |
|---|---|---|---|
| Execute bounded | Low-risk system actions inside well-defined rules, with audit logs and rollback paths. | Human sets policy, reviews exceptions, and monitors metrics. | Use only where failure impact is limited and verification is strong |
| Restricted / support-only | Eligibility, adverse decisions, care-impacting actions, disciplinary actions, or other high-impact determinations. | Human remains the decision-maker; AI can structure evidence but not silently decide. | Default posture for consequential public-sector workflows |

Federal policy basis: [1] M-25-21 requires proportionate risk management for high-impact AI and emphasizes trustworthy, accountable use; [2] M-25-22 emphasizes documentation, transparency, and accessibility; [3] NIST AI RMF supports risk-tiered implementation.

## 5. Guardrails and trust architecture

A credible public-sector AI strategy must promise speed and discipline at the same time. The most useful guardrails are the ones that let AI accelerate safe work while forcing attention to the points where rights, safety, mission integrity, or public trust are genuinely at stake. They should be built into the delivery method—not bolted on after a pilot looks impressive.

Government teams should assume that AI can be confidently wrong, incomplete, outdated, or misaligned to local operating reality. The answer is not to ban it from everything. The answer is to engineer bounded autonomy, independent verification, provenance, and clear human accountability.

| Guardrail | Why it matters | Minimum practice |
|---|---|---|
| Should-we gate | Not every use case that is technically feasible is appropriate for public-sector automation. | Screen use cases for rights impact, public trust, reversibility, and mission criticality before building. |
| Authoritative context packs | Ungrounded models produce plausible but generic output. | Require approved-source retrieval and current context for any material requirement, recommendation, or draft artifact. |
| Data provenance and synthetic-data controls | Teams need to know what is authoritative, what is derived, and what may be synthetic or stale. | Track source provenance and favor authoritative human-created or system-of-record content for consequential work. |
| Human accountability at the right loop | People, not models, remain responsible for agency decisions and public communications. | Assign named human owners for approval, exception handling, and risk acceptance. |
| Separation of planning and action | A system that both decides and executes can compound its own errors. | Use independent checkpoints between recommendation and execution in consequential workflows. |
| Verify the state, not the statement | Agent-reported completion is not evidence. | Validate against logs, records, transaction states, or external system checks before counting work as done. |

| | | |
|---|---|---|
| Non-public data protections | Agency data should not leak into public or commercial model training without approval. | Ensure terms, architecture, and contracts prohibit such training from absent explicit agency consent. [2] |
| Accessibility, plain language, and explainability | Public-facing automation that is inaccessible or opaque erodes trust and increases downstream burden. | Build Section 508, plain language, and explanation requirements into prompts, templates, tests, and reviews. |
| Risk-tiered monitoring and discontinuation criteria | High-impact AI that underperforms should not simply continue by inertia. | Define performance thresholds, escalation rules, pause criteria, and discontinuation paths for sensitive use cases. [1] |
| Full traceability | Without traceability, reuse and oversight break down. | Preserve links among need, source context, requirements, generated artifacts, decisions, tests, approvals, and operational feedback. |

Guardrail basis: [1] M-25-21; [2] M-25-22; [3] NIST AI RMF / Generative AI Profile; [4][5] federal CX and digital-experience requirements. The specific operating practices here are product-agnostic design recommendations.

## 6. Better outcomes for the public, staff, and mission owners

The point of contextual AI is not novelty. It is better outcomes. For federal departments and agencies, that means clearer public and workforce journeys, faster response and processing, fewer avoidable errors, more consistent information across channels, lower backlog friction, and systems that are easier to change without breaking everything around them. For staff, it means less swivel-chair work, better access to knowledge, faster artifact creation, clearer traceability, and more time for judgment and service.

| Who benefits | Better outcome | How contextual AI contributes | Illustrative measures |
|---|---|---|---|
| The public / customers | Faster, clearer, more consistent service | Grounded draft notices, better intake, better triage, easier self-service, and more coherent multichannel experiences | Cycle time, error rate, accessibility defects, satisfaction |
| Frontline staff | Less search burden and manual re-entry | Retrieval support, draft responses, next-step suggestions, and structured work products built from real context | Handle time, first-contact resolution, hours redirected |
| Program leaders | Better portfolio visibility and lower rework | More structured intake, reuse discovery, traceability, and evidence-based release decisions | Backlog aging, reuse rate, release predictability |
| Technology and risk leaders | More maintainable and scalable solutions | Generation against approved patterns, better control evidence, clearer option analysis, and stronger audit trails | Defect escape rate, approval latency, findings closed |

Public success inputs relevant to these outcomes include reusable shared services, service-centered modernization, customer experience programs, digital-form modernization, and high-visibility delivery efforts such as identity, filing, and inspection workflows. [11][12]

## 7. How this generalizes across the U.S. Government

This approach is not limited to one program name, administration, or acquisition vehicle. Across the U.S. Government, the same model can support benefits administration, health and human services, tax administration, contact centers, identity services, grants, inspections, enterprise shared services, and internal mission support. The same operating logic is relevant wherever agencies manage high-volume service delivery, casework, eligibility decisions, regulated workflows, or complex legacy environments. Government-wide AI and digital-service guidance emphasize reuse, coordination, and enterprise maturity. [1][13]

| Mission pattern | Representative federal contexts | Illustrative agencies / environments |
|---|---|---|
| Benefits and casework | Benefits, eligibility, case management, appeals, relief, and survivor assistance | SSA, FEMA, USDA, HHS, Labor, and similar assistance programs |
| Health and public service operations | Public health, scheduling, logistics, service operations, and call centers | HHS, IHS, CDC, FEMA, and other complex service operations |
| Citizen contact and self-service | Contact centers, identity, digital forms, notices, web and mobile interactions | IRS, SSA, USCIS, GSA, FEMA, and other high-volume public contact organizations |
| Grants, inspections, and regulatory review | Permitting, inspections, investigations, evidence assembly, status tracking | HUD, EPA, Labor, Commerce, DOT, USDA, and similar review-heavy environments |
| Mission-support shared services | Acquisition, HR, finance, IT support, training, records, security, compliance | Treasury, GSA, DHS, DOJ, OPM, and cross-agency shared-service environments |

## 8. 12-month adoption path

The most credible path is phased. Start where the context is available, the risk is bounded, and the economics are clear. Prove that contextual AI improves throughput, quality, and reuse in a few lanes before expanding into broader modernization or more sensitive decisions.

| Window | What to do | What to prove |
|---|---|---|
| 0–90 days | Stand up principles, use-case intake, risk tiers, governance roles, and context-pack patterns. Prioritize a few low- or medium-risk pilot lanes. | That the agency can assemble authoritative context reliably and route work through clear decision rights. |

| | | |
|---|---|---|
| 90–180 days | Pilot AI-assisted intake, Sprint Zero, fit-gap analysis, testing, knowledge support, or support-content creation. Instrument the work heavily. | Measured cycle-time reduction, lower rework, usable traceability, and acceptable human-review burden. |
| 180–365 days | Expand to additional lanes, institutionalize reusable templates, connect to enterprise architecture and approval workflows, and scale training. | Repeatability across teams, stronger reuse rates, more predictable release readiness, and clearer benefit to end users. |

## Illustrative KPI Targets

The following illustrative KPI targets are discussion-grade ranges that can be used to define pilots, compare against a recent baseline, and decide what should scale.

| Metric | Illustrative 12-Month Target |
|---|---|
| Intake to approved backlog cycle time | 40-60% faster |
| Approved backlog to release candidate | 25-40% faster |
| Rework caused by requirements gaps, missing context, or late control findings | 25-50% lower |
| Quality and decision latency (defect escape, late accessibility defects, approval latency) | 20-35% lower defect escape;<br>40-60% fewer late accessibility defects;<br>30-50% faster release or risk approvals |
| Reuse rate for templates, patterns, prompts, workflows, and controls across teams | 2x-3x higher |
| Hours redirected from manual coordination, drafting, or repetitive processing | 10-20% of team capacity shifted to higher-value work |
| Knowledge-support performance (search success, handle time, first-contact resolution, staff confidence) | Search success +10-20 points; handle time 10-20% lower;<br>first-contact resolution +5-10 points;<br>staff confidence +15-25 points |
| End-user outcomes (service cycle time, clear-notice rate, digital completion, satisfaction) | Service cycle time 10-25% faster;<br>clear-notice rate +10-20 points;<br>digital completion +5-15 points;<br>satisfaction +5-10 points |
| Audit and oversight (trace completeness, control coverage, issues closed) | 90%+ trace completeness for pilot artifacts;<br>95%+ mapped control coverage;<br>20-30% faster issue closure |

## Conclusion

The predominant AI position should not be "AI everywhere." It is contextual AI where it changes the economics and quality of delivery: from intake, requirements, decomposition, modernization decision-making, to testing, knowledge support, operations, and reuse. That is how government agencies can move faster without becoming less trustworthy.

For federal departments and agencies, this approach supports better outcomes for the public and the staff who serve them. Across the wider federal enterprise, it offers a product-agnostic operating model that can be adapted to similar missions and constraints. The objective is not to remove judgment. It is to remove wasteful translation loops, improve evidence and maintainability, and allow people to spend more time on the parts of public service that require human responsibility.

As federal departments and agencies continue to modernize the digital enterprise, the imperative is not simply to adopt artificial intelligence and apply it to existing methods and processes, but to apply it in ways that are mission-aligned, governed, transparent, and accountable. When used responsibly across the lifecycle of digital strategy, requirements, delivery, operations, and continuous improvement, AI can help agencies improve timeliness, strengthen decision support, reduce avoidable burden on the workforce, and deliver better services to the public. The agencies that realize the greatest value will be those that pair innovation with discipline, combining contextual data, human judgment, paired learning models, and sound governance to achieve measurable outcomes. For organizations seeking to better understand how AI can support lifecycle management of the digital enterprise in a practical and responsible way, **Elevate Federal** welcomes the opportunity for further discussion. **Contact Elevate Federal** at **info@elevatefederal.com**

## Selected public sources

The following sources were used by Elevate Federal to anchor the public-record observations in this paper and provide a baseline for tailoring the approach to specific federal department and agency use cases.

**[1]** OMB Memorandum M-25-21, Accelerating Federal Use of AI through Innovation, Governance, and Public Trust (2025).
https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-21-Accelerating-Federal-Use-of-AI-through-Innovation-Governance-and-Public-Trust.pdf

**[2]** OMB Memorandum M-25-22, Driving Efficient Acquisition of Artificial Intelligence in Government (2025).
https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-22-Driving-Efficient-Acquisition-of-Artificial-Intelligence-in-Government.pdf

**[3]** NIST AI Risk Management Framework 1.0 and NIST-AI-600-1, Generative AI Profile.
https://www.nist.gov/itl/ai-risk-management-framework

https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

**[4]** Executive Order 14058 and Federal Customer Experience guidance.
https://www.performance.gov/cx/executive-order/

https://www.performance.gov/cx/

**[5]** 21st Century IDEA implementation resources and GSA Digital Experience guidance.
https://www.gsa.gov/technology/government-it-initiatives/digital-experience

**[6]** GAO-25-107795, Information Technology: Agencies Need to Plan for Modernizing Critical Decades-Old Legacy Systems (Jul. 2025).
https://www.gao.gov/products/gao-25-107795

**[7]** GAO-25-106908, Mission-Critical Information Technology: Agencies Are Monitoring Selected Acquisitions for Cybersecurity and Privacy Risks (Mar. 2025).
https://www.gao.gov/products/gao-25-106908

**[8]** GAO-25-107852, High-Risk Series: Critical Actions Needed to Urgently Address IT Acquisition and Management Challenges (Jan. 2025).
https://www.gao.gov/products/gao-25-107852

**[9]** GAO-24-106764, Digital Experience: Agency Compliance with Statutory Requirements (Sep. 2024).
https://www.gao.gov/products/gao-24-106764

**[10]** GAO-26-107859, Artificial Intelligence Acquisitions: Agencies Should Collect and Apply Lessons Learned to Improve Future Procurements (Apr. 2026).
https://www.gao.gov/products/gao-26-107859

**[11]** Technology Modernization Fund FY24 Annual Report, TMF impact materials, and USDA Specialty Crops Inspection case study.
https://tmf.cio.gov/impact/
https://tmf.cio.gov/files/2024/12/TMF-FY24-Annual-Report.pdf
https://tmf.cio.gov/usda-case-study/

**[12]** Federal digital service execution inputs: Login.gov, IRS Direct File, and the Digital Services Playbook.
https://www.gsa.gov/blog/2024/08/26/gsas-logingov-expands-services-into-states
https://www.irs.gov/pub/irs-pdf/p5969.pdf
https://playbook.usds.gov/

**[13]** America's AI Action Plan (Jul. 2025) for cross-government adoption, coordination, and enterprise AI maturity.
https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf