Evolver

AI-Powered Information Technology (IT) Consolidation and Modernization:

U.S. Federal Government Challenges, Tools, & Best Practices

By: Gregory A. Garrett, Rahul Johri, & Dr. Brian McElyea

Executive Summary

The United States Federal government is currently navigating a period of profound business and technological uncertainty. This uncertainty stems from the convergence of aging legacy information technology (IT) infrastructure, a shrinking technical and acquisition workforce, a major overhaul of the cumbersome and rigid Federal Acquisition Regulatory (FAR) system, budgetary pressure, and exponential growth in data volume.

Together, these forces have created critical operational gaps that traditional IT service delivery models can no longer bridge effectively. Artificial Intelligence (AI) enters this gap, not merely as a collection of automation tools, but as a transformative architectural paradigm that promises government operations from reactive administration to predictive, autonomous service delivery.

This whitepaper, "AI-Powered Information Technology Consolidation and Modernization," delivers a strategic analysis tailored for U.S. Federal government agencies. It explores the mechanisms, challenges, and best practices essential for navigating this AI-driven transformation. The urgency for this shift is driven by stark realities: U.S. Federal agencies face massive backlogs, from Veterans Affairs (VA) benefit claims, Social Security Administration benefits processing, to Internal Revenue Service (IRS) correspondence, that human labor alone cannot resolve within acceptable timeframes. [1] At the same time, the technological foundation of the U.S. government remains brittle, with critical systems that still depend on COBOL codebases dating back to the 1960s, creating severe risks to operational continuity of government operations. [2]

AI-powered IT services offer a twofold approach: immediate operational relief through tools like Generative AI (GenAI) chatbots and Robotic Process Automation (RPA), and long-term structural remediation through AI-assisted code refactoring and AI IT Ops (Artificial Intelligence for IT Operations). The market for these services is expanding rapidly, with vendors now assessed not just on technical capability, but on their ability to navigate the complex regulatory landscape defined by Executive Order 14110 and the Office of Management and Budget (OMB) Memorandum M-24-10. [3]

This paper examines the transition from "basic automation" to "full autonomy," where agentic AI systems are capable of independent decision-making within bounded parameters.[4] It addresses the critical business challenges, including the data readiness gap and the "silver tsunami" of Federal employee retirements expected over the next 3-5 years, and outlines key tools and best practices, such as FedRAMP High authorization, Zero Trust architectures, and the NIST AI Risk Management Framework that underpin successful federal information systems implementation. Our intent is to deliver a practical roadmap for government and industry partners seeking to leverage AI to facilitate the cost-effective consolidation and modernization of U.S. Federal enterprise IT.

Current AI systems, including advanced large language models (LLMs), are not sentient; i.e., they do not possess consciousness, self-awareness, or subjective experience. Federal policy and scientific consensus frame AI as pattern-based automation, not cognition. Executive Order 14110 and the NIST AI RMF emphasize risk-managed autonomy with human oversight, avoiding speculative claims about sentience. This distinction is critical for compliance and public trust.

AI-Powered IT Modernization for U.S. Federal Government

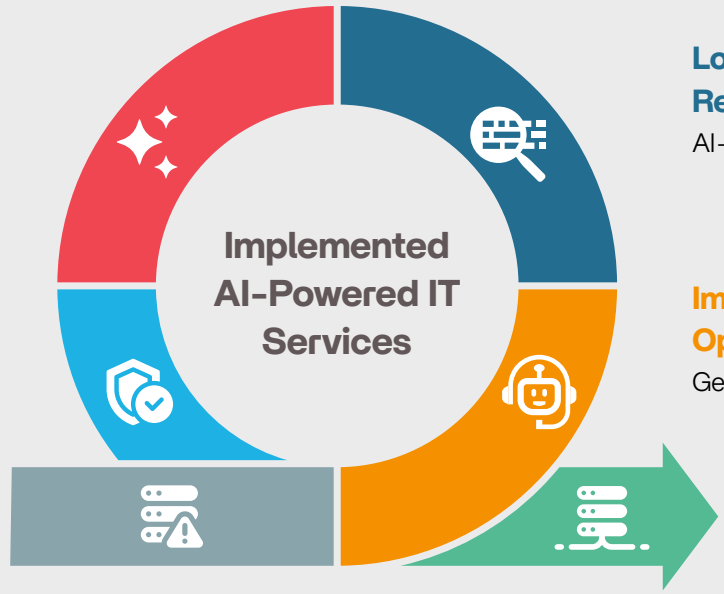

Figure 1: AI Modernization Roadmap – From Basic Automation to Autonomous IT Operations.

Key AI Implementation Challenges

Federated IT Governance Structure

Numerous U.S. Federal government agencies still operate within heavily federated IT environments that include multiple Offices of Information Technology. This creates layers of Chief Information Officers (CIOs), Chief Information Security Officers (CISOs), and Chief Technology Officers (CTOs) across headquarters, centers, institutes, and/or commands, each with their own IT policies, systems, and staffing. These fragmented structures often become a barrier to driving innovation, making it more difficult to achieve system interoperability. The fragmentation slows down modernization efforts, creates inconsistent technology practices, and increases operations and maintenance costs. The result is a governance model that struggles to keep pace with the speed, scale, and integration demands of AI and mission needs.

Legacy IT Infrastructure and Technical Debt

The most pervasive barrier to AI adoption in the U.S. Federal government is the entrenched prevalence of legacy IT systems. A Government Accountability Office (GAO) review identified numerous critical systems ranging from 8 to 51 years old, many written in obsolete languages like COBOL and Assembly. [5] These systems were designed in an era of batch processing and on-premises mainframes. They are inherently hostile to modern AI pipelines that require real-time data ingestion, API-first connectivity, and elastic cloud compute.[6]

The risk posed by these legacy IT systems is existential. As the pool of engineers fluent in COBOL retires, the knowledge required to maintain these systems evaporates. This “legacy code challenge” creates a paradox: agencies

need AI to modernize their systems, but their systems are often too antiquated to support AI integration. For example, integrating a modern Large Language Model (LLM) to query a mainframe database requires building complex API layers that the original system architecture never anticipated, often leading to performance bottlenecks or data integrity issues.[7]

Furthermore, legacy IT systems often lack the comprehensive digital instrumentation required for AI Observability. AIOps platforms, which use machine learning to predict IT outages, require high-fidelity telemetry data (logs, metrics, traces).[8] Legacy mainframes often output opaque or unstructured logs that modern AIOps tools struggle to ingest without significant ETL (Extract, Transform, Load) processes. This limits the ability of agencies to apply predictive maintenance to their most critical and fragile assets, leaving them reactive to outages rather than proactive in prevention.[9]

Data Readiness and Sovereignty

AI is only as effective as the data it is trained on or has access to. For most federal agencies, data readiness is a profound challenge. Government data is frequently siloed across disparate bureaus, centers, and institutes, while stored in incompatible formats, and laden with quality issues such as duplication, incompleteness, and lack of standardized metadata.[10] "Dirty data" can lead to hallucinations in GenAI models or inaccurate predictions in decision-support systems, creating unacceptable risks in mission-critical contexts.

The challenge extends to data sovereignty and residency. U.S. Federal agencies, particularly those dealing with national security or sensitive citizen data (Controlled Unclassified Information, or CUI), have strict requirements regarding where data can be stored and processed. The "sovereign cloud" model is increasingly relevant, where AI workloads must run in physically isolated environments "enclaves" or within specific geographic boundaries to meet regulatory mandates.[11] Agencies must navigate the complexities of FedRAMP authorization, ensuring that any commercial AI service provider meets the rigorous cybersecurity controls of FedRAMP Moderate or High baselines depending on the data's impact level.[12]

Moreover, the integration of commercial "black box" AI models raises questions about data leakage. If an agency uses a public LLM for summarization or code generation, there is a risk that sensitive government data could be ingested into the model training set, potentially exposing it to external actors. OMB M-24-10 explicitly requires contracts to prohibit the use of non-public government data for training commercial algorithms without explicit consent, adding a layer of contractual complexity to IT service procurements.[13]

The Federal AI Workforce Gap

The U.S. Federal government and industry face a critical shortage of personnel with the specialized skills necessary to acquire, deploy, and manage AI systems. This "AI talent gap" is exacerbated by intense competition from the private sector, where compensation for AI engineers and data scientists significantly outpaces federal salary scales.[14] The 2024 GAO report highlighted that agencies like the Department of Commerce, DoD, and NASA struggle to attract and develop individuals with generative AI expertise.[15]

This shortage is not limited to technical roles; it extends to the federal acquisition workforce. Contracting officers often lack the training to evaluate complex AI proposals or to structure performance-based contracts for non-deterministic technologies. The AI Training Act was signed into law to address this, mandating AI training for the federal acquisition workforce. However, the implementation of these programs is an ongoing process.[16]

The impending workforce cliff (the retirement of a significant portion of the federal employee workforce, including both federal acquisition and technical personnel) compounds this issue. As experienced staff leave, they take institutional knowledge with them. While AI knowledge management systems offer a partial solution by capturing and indexing this information, the immediate effect is a "brain drain" that leaves agencies vulnerable during the transition to AI-enabled operations.[17]

Agency Challenge	Description	AI Impact
Legacy Code	Reliance on COBOL/Mainframes (e.g., IRS, Treasury).	Necessitates API layers or AI-refactoring; increases modernization cost.
Data Hygiene	Siloed, unstructured, or duplicated data.	Leads to model hallucinations; reduces predictive accuracy.
Talent Shortage	Lack of AI engineers and data scientists.	Delays deployment; increases reliance on expensive contractors.
Acquisition	Slow procurement cycles (18-24 months).	Tech becomes obsolete before award; difficulty vetting AI vendors.
Sovereignty	Strict requirements for CUI/National Security data.	Limits use of public commercial AI models; mandates GovCloud.

Table 1: Agency Challenges and Potential Impact

Cybersecurity, Trust, and Adversarial AI

Cybersecurity in the AI era extends beyond traditional perimeter defense. Agencies must now contend with "Adversarial AI" cyber-attacks specifically designed to manipulate AI systems. This includes "data poisoning," where adversaries inject malicious data into training sets to corrupt the model's behavior, and "model inversion," where cyber-attackers query a model to reconstruct the sensitive data it was trained on.[18]

The Department of Homeland Security (DHS) and other agencies have identified "High Impact" AI use cases, those that impact rights or safety, as requiring enhanced risk management.[19] Ensuring that AI systems are explainable, fair, and free from bias is a massive challenge. A "black box" algorithm cannot simply be deployed to adjudicate benefits or screen cargo; agencies must be able to audit why the AI made a specific decision. The NIST AI Risk Management Framework (RMF) provides guidance for this (Govern, Map, Measure, Manage), but operationalizing these abstract principles into IT service contracts remains a hurdle for many agencies.[20]

Furthermore, the concept of "Shadow AI" is emerging as a significant threat. Employees, driven by the desire for efficiency, may unofficially use unvetted public AI tools (like ChatGPT) for official work, bypassing cybersecurity protocols and potentially leaking CUI.[21] IT service providers must therefore not only deploy approved AI tools but also implement CASB (Cloud Access Security Broker) solutions to detect and block unauthorized AI usage.

AI Tools & Best Practices

To navigate the challenges outlined above, U.S. federal agencies and their industry partners are deploying a suite of advanced technologies and adhering to rigorous best practices. The focus is moving quickly from conducting experimental AI "pilots" to scalable, AI-powered enterprise-grade solutions that are secure, observable, compliant, and cost-effective. The AI ecosystem in U.S. federal IT is diverse, encompassing everything from predictive infrastructure monitoring to citizen-facing conversational agents. These are all designed to help facilitate cost-effective IT consolidation and modernization to support the needs of the nation.

Generative AI and Large Language Models (LLMs) Tools

Generative AI has seen an explosion in adoption, with reported use cases in federal agencies increasing nine-fold between 2023 and 2024.[22] The deployment of these models is characterized by specific architectural choices designed to mitigate risk.

- **Retrieval-Augmented Generation (RAG):** Agencies are moving away from training models from scratch due to cost and data privacy concerns. Instead, they are adopting RAG architectures. In a RAG system, a generative model is connected to a trusted, internal knowledge base (e.g., policy documents, technical manuals). When a user asks a question, the system retrieves relevant documents and uses the LLM to summarize an answer based only on that trusted data. This minimizes hallucinations and ensures answers are grounded in official policy. This approach is being used by the U.S. State Department in its "StateChat" internal tool.[23]
- **Code Refactoring Agents:** To address the COBOL legacy issue, agencies are employing GenAI tools trained on mainframe languages to automatically document, analyze, and refactor legacy code into modern languages like Java or Python. Tools like "GenWizard" and platforms from GitLab are being used to accelerate this modernization, reducing the dependency on the dwindling cohort of COBOL programmers.[24]
- **Agentic AI:** The next frontier is "Agentic AI," systems that do not just answer questions but take autonomous actions to complete workflows. For example, an AI agent could independently identify a server outage, open a ticket in ServiceNow, run a diagnostic script, and apply a patch, all without human intervention.

AI IT Ops and Observability Tools

As IT environments become more hybrid and complex, human operators can no longer manually monitor system health. AI IT Ops (Artificial Intelligence for IT Operations) platforms are becoming standard.

- **Predictive Remediation:** Platforms like Dynatrace and ScienceLogic use Causal AI to analyze causal relationships in infrastructure data. Instead of just alerting on a high CPU spike, these tools trace the spike back to a specific bad code deploy or database lock and can automatically trigger a rollback.[25]
- **Network Operations Center (NOC) Virtualization:** Agencies are "virtualizing" their NOCs using AI to consolidate tools and reduce alert fatigue. ScienceLogic's deployment at a large U.S. media provider (analogous to a large agency) reduced incident response time by 90% by automating ticket triage and remediation.[26]
- **Full-Stack Observability:** SolarWinds and similar vendors offer "full-stack" visibility that integrates network, database, and application monitoring. This is crucial for "High Impact" systems where uptime is a matter of national security or public safety.[27]

Robotic Process Automation (RPA) Tools

RPA remains the "digital duct tape" of U.S. federal IT, bridging the gap between modern web interfaces and legacy backend systems that lack Application Programming Interfaces (APIs).

- **Intelligent Automation:** Modern RPA is evolving into "Intelligent Automation" by integrating with AI. For instance, a bot can extract data from a scanned PDF form (using AI Optical Character Recognition), classify the document (using Natural Language Processing), and then input the data into a legacy mainframe (using RPA
- **Use Cases:** The General Services Administration (GSA) uses RPA/AI to review contract documents, significantly reducing processing time. [28] UiPath has extensively documented cases where bots handled surge workloads for unemployment claims during the pandemic, processing transactions in hours that would have taken humans hundreds of hours.[29]

AI-Powered Cybersecurity Tools

With the rise of automated cyber-attacks, agencies are deploying AI-driven defense mechanisms.

- **Endpoint Detection and Response (EDR):** Tools like CrowdStrike's Charlotte AI use generative AI to allow security analysts to query threat data using natural language (e.g., "Show me all endpoints that communicated with this malicious IP in the last 24 hours"). This democratizes threat hunting and speeds up response times.[30]
- **Risk-based Vulnerability Management (RBVM):** Tools like Evolver's GUARDIAN solution which combines Evolver's proven cybersecurity governance, risk, and compliance (GRC) and risk management framework (RMF) services with Nucleus Security's FedRAMP Moderate authorized AI-powered automated unified vulnerability and exposure management software platform, which have proven to accelerate the cyber risk reduction and prioritization process by 80%.
- **Enterprise Cyber Threat Recognition and Analysis:** Tools like Evolver's SPECTRA solution which combines a self-supervised AI overlay, powered by MixMode AI with a 24/7 Security Operations Center cybersecurity services provided by Evolver. The Evolver SPECTRA solution has led to reducing the Mean-Time-To-Detect (MTTD) and Mean-Time-To-Remediate (MTTR) by up to 70% with fewer non-actionable security alerts.
- **Zero Trust Architecture (ZTA) Integration:** AI is essential for Zero Trust architectures, which require continuous authentication and behavioral monitoring. AI models analyze user behavior baselines to detect anomalies (e.g., a user downloading unusual volumes of data) that might indicate an insider threat or compromised credential.
- **Cybersecurity Supply Chain Risk Management (C-SCRM):** Obtaining cybersecurity self-assessments from government contractors/subcontractors/suppliers is inadequate to ensure real cybersecurity. Using time consuming and expensive third-party provided cyber compliance assessments of suppliers, based upon static government cyber compliance frameworks is a step in the right direction, but still largely ineffective in a world of evolving cyber threats. Tools like Evolver's CLEAR solution which pairs Evolver's cyber RMF services with WhiteHawk's proven effective cyber-risk analytics platform, which delivers independent, fast, zero-intrusion supplier cyber risk assessments mapped to customer selected cyber risk frameworks – takes C-SCRM to a new level of performance.

AI Implementation Best Practices

Successful AI implementation in government requires adherence to a rigorous set of management and technical practices.

The NIST AI Risk Management Framework (RMF)

The NIST AI RMF (NIST SP 800-37 Rev. 2) is the gold standard for U.S. Federal AI governance. It is organized around four core functions:

1. **Govern:** Establish policies, accountability structures, and a culture of risk management. This involves designating a CAIO and defining risk tolerance.[31]
2. **Map:** Contextualize the risks. Agencies must document the intended purpose of the AI, its limitations, and the potential impacts on specific populations. This is the "inventory" phase mandated by OMB.[32]
3. **Measure:** Quantify risks using metrics for accuracy, robustness, and fairness. This involves rigorous testing and evaluation (T&E) before deployment.[33]
4. **Manage:** Prioritize and act on risks. This includes implementing controls, monitoring for "model drift" (where AI performance degrades over time), and having a "kill switch" for systems that behave erratically.[34]

Compliance with OMB M-24-10

OMB Memorandum M-24-10 establishes binding requirements for AI acquisition.

- **Rights-Impacting vs. Safety-Impacting:** Agencies must classify their AI use cases. If a system impacts

rights (e.g., hiring, benefits, law enforcement) or safety (e.g., critical infrastructure, medical devices), it triggers "Minimum Risk Management Practices."

- **Mandatory Safeguards:** These practices include performing an AI impact assessment, independent evaluation, and ongoing monitoring. Agencies must also provide a mechanism for the public to opt-out or appeal AI decisions.[35]
- **Transparency:** Agencies must publish annual inventories of their AI use cases, identifying which are rights- or safety-impacting. DHS, for example, published a detailed inventory in 2024 listing 158 active use cases.[36]

Federal agent deployments must align with:

- EO 14110
- OMB M-24-10
- NIST AI RMF
- FedRAMP High for sensitive data
- Zero Trust principles
- Human-in-the-loop for high-stakes decisions

FedRAMP High Authorization

For AI systems handling sensitive data, FedRAMP High authorization is non-negotiable. This requires:

- **421+ Controls:** A comprehensive set of security controls covering access control, encryption, and incident response.
- **Physical Isolation:** Often requires data to be stored in "GovCloud" regions that are physically separated from commercial public cloud infrastructure and staffed only by U.S. citizens.[37]
- **Continuous Monitoring:** Authorization is not a one-time event; providers must submit monthly vulnerability scans and undergo annual third-party assessments (3PAO).[38]

Human-in-the-Loop (HITL)

For high-stakes decisions, complete autonomy is rarely acceptable. Best practice dictates a "Human-in-the-Loop" or "Human-on-the-Loop" design, where the AI provides a recommendation or draft, but a human officer makes the final adjudication. This is critical for complying with the "administrative due process" requirements inherent in government service.[39]

U.S. Federal Government Sector-Specific: AI Best Practices

The application of AI varies significantly across different sectors of the federal government, driven by unique mission requirements.

Civilian Agencies: Citizen Services, Consolidation, & Cost Effectiveness

- **Voice Bots & AI Chatbots:** The IRS deployed voice bots on toll-free lines to handle simple queries like payment plan setup and transcript requests. These bots, powered by AI, have answered millions of calls, reducing wait times for taxpayers with complex issues who need to speak to a human agent.[40]
- **Shared AI-powered IT & Cybersecurity Tools and Services**: For years the U.S. Department of Homeland Security (DHS) has offered and provided a wide range of IT and cybersecurity support services to numerous U.S. Federal Civilian Executive Branch (FCEB) agencies via the shared-services program and via the DHS Continuous Diagnostics and Mitigation (CDM) program. Now those programs are being expanded to include AI-powered IT and cybersecurity tools.
- **Enforcement:** AI is used to analyze complex partnership returns and identify tax evasion patterns that would be invisible to manual auditors. This high-impact use case helps close the tax gap.
- **Claims Processing:** The U.S. Department of Veteran Affairs (VA) utilizes AI to classify and extract data from millions of pages of medical evidence submitted by veterans. This automation has been credited with reducing the claims backlog by 57%, allowing adjudicators to focus on decision-making rather than data entry.[41]

- **Healthcare:** In the medical domain, the VA is piloting GenAI to summarize patient history for clinicians and using predictive AI to forecast appointment demand, optimizing staffing levels.[42]
- **StateChat:** The U.S. State Department launched "StateChat," an internal, secure GenAI chatbot. It allows diplomats to summarize cables, draft speeches, and translate documents using "Sensitive But Unclassified" (SBU) data. This tool is designed to save thousands of work hours while keeping data within the Department's secure boundary.[43]

Department of War & Department of Homeland Security: Speed, Security, and Autonomy

The Department of War (DoW)'s focus is on "decision superiority" and contested logistics.

- **Project Maven:** A flagship AI initiative, Maven uses computer vision to analyze drone footage and satellite imagery to identify targets and threats autonomously. This dramatically speeds up the "OODA loop" (Observe, Orient, Decide, Act).[44]
- **Contested Logistics:** AI is used to predict maintenance failures on vehicles and aircraft ("predictive maintenance") and to optimize supply chains in real-time, ensuring that troops have ammunition and fuel even when supply lines are disrupted.[45]
- **Generative AI for Doctrine:** Leidos and other partners are working with the Army to turn static doctrine documents into interactive GenAI tools, allowing soldiers to query field manuals via voice or text in the heat of operation.[46]

Department of Homeland Security (DHS) uses AI for border security, cyber threat detection, and much more.

- **Cargo Screening:** AI algorithms analyze X-ray images of cargo trucks and airport luggage to detect anomalies (weapons, drugs) with higher accuracy and speed than human screeners.[47]
- **Inventory Transparency:** DHS has been a leader in transparency, publishing a detailed inventory of its AI systems, including those used for facial recognition at borders, and detailing the risk mitigations in place.[48]

Future Trends in Federal AI Modernization

As U.S. federal agencies progress from basic automation toward autonomous IT operations, several emerging trends will shape the next decade of modernization:

- **Multi-Agent Systems and Agentic AI at Scale:** Agentic AI will evolve beyond single-task autonomy into orchestrated multi-agent ecosystems. These systems will collaborate to manage complex workflows, such as predictive maintenance, cybersecurity response, and supply chain optimization, while adhering to risk-managed autonomy principles outlined in EO 14110 and NIST AI RMF.
- **AI + Quantum Computing Integration & PQC Security:** The convergence of AI and quantum computing will accelerate optimization tasks, cryptographic resilience, and large-scale simulations. U.S. government agencies must prepare for post-quantum cryptography (PQC) standards to secure AI-driven systems against future threats leveraging capabilities like Evolver's SHIELD solution which combines Quantum Xchange's crypto agility Phio-TX capabilities, which are NIST approved and FIPS validated, with Evolver's cybersecurity services to deliver cost-effective enterprise quantum security at scale.
- **Synthetic Data for Privacy-Preserving AI:** To address strict data sovereignty and privacy mandates, agencies are increasingly adopting synthetic datasets for AI model training. This strategy minimizes reliance on sensitive citizen information while preserving model accuracy and ensuring compliance with OMB M-24-10 safeguards. Additionally, synthetic data strengthens defenses against adversarial attacks, such as attempts to reverse-engineer real data from AI models.
- **Automated AI Governance and Policy Engines:** Future compliance will move from static frameworks to dynamic enforcement. AI-driven policy engines will automatically apply OMB and NIST guardrails, monitor model drift, and trigger "kill switches" for erratic behavior, creating real-time governance for high-impact systems.

- **Automated AI Governance and Policy Engines:** Future compliance will move from static frameworks to dynamic enforcement. AI-driven policy engines will automatically apply OMB and NIST guardrails, monitor model drift, and trigger "kill switches" for erratic behavior, creating real-time governance for high-impact systems.
- **Edge AI for Tactical and Mission-Critical Operations:** Defense and homeland security agencies will deploy AI at the tactical edge, enabling real-time decision-making in contested environments. By processing data locally, Edge AI supports autonomous logistics, rapid threat detection, and secure communications, thereby reducing reliance on centralized cloud infrastructure and keeping sensitive information compartmentalized.
- **Humanoid Robotics for Human-AI Collaboration:** Future federal modernization may include humanoid robotic systems designed for mission-critical environments, such as logistics support, hazardous material handling, and disaster response. These robots will integrate advanced AI for perception, navigation, and task execution, enabling seamless collaboration with human operators while adhering to safety and compliance standards.

Summary and Conclusion

The integration of AI-Powered information technology consolidation and modernization into the U.S. Federal government is no longer a theoretical aspiration; it is an operational imperative. Faced with the multiple pressures of a shrinking workforce, need for reduced budgets, and an exploding demand for services, agencies are rapidly pivoting from small AI pilot programs to large AI enterprise-scale adoption. The evidence is clear: from the IRS's voice bots handling millions of calls to the VA's 57% reduction in claims backlogs, AI is delivering tangible, high-impact results.

However, this transformation is fraught with complexity. The "legacy anchor" of COBOL-based infrastructure remains a significant drag on modernization, requiring innovative approaches like Evolver's new TRANSFORM solution, leveraging Rhino.ai's Universal Application Notation (UAN) capability, with low-code/no-code software platforms like Appian to rapidly modernize legacy software in 70% less time. The talent gap poses a severe risk to the long-term sustainability of these systems, necessitating aggressive upskilling initiatives like the Federal AI Training Act and the U.S. Digital Corps fellowship programs.

For U.S. government contractors, the message is unequivocal. Success in this new era requires more than just technical prowess. It demands a deep understanding of the regulatory landscape, specifically the "high impact" designations of OMB M-24-10 and the risk management protocols of NIST. Vendors must be prepared to offer sovereign AI solutions that guarantee data residency and security, often at the FedRAMP Moderate to High level. Government contractors must shift from selling highly customized software products to delivering cost-effective integrated solutions, leveraging leading commercial Software-as-a-Service (SaaS) with transparent, explainable, and trustworthy AI ecosystems.

The future of U.S. federal IT is agentic, predictive, and autonomous. As agencies move toward this future, the partnership between government and industry will evolve from a transactional relationship to a strategic alliance focused on co-creating value. The roadmap is drawn; the challenge now is execution. By adhering to the principles of trust, cybersecurity, and human-centric design, the U.S. federal government can harness the power of AI to build a more responsive, efficient, and resilient public service for the 21st century and beyond.

Works Cited

"VA Reduces Backlog of Veterans Waiting for VA Benefits by 57%." U.S. Department of Veterans Affairs, 18 Nov. 2025. Accessed 12 Dec. 2025. [1, 2, 41, 42]

"Fact Sheet: OMB Issues Guidance to Advance the Responsible Acquisition of AI in Government." The White House, 3 Oct. 2024. Accessed 12 Dec. 2025. [3, 13, 19, 35, 36]

"How AI Can Fix Government's Legacy Code Problem." GitLab, 29 July 2025. Accessed 12 Dec. 2025. [4, 5, 6, 24]

"The Role of AI in Legacy System Modernization." Webelight Solutions. Accessed 12 Dec. 2025. [7, 8, 9, 10, 25, 40]

"AI-Powered Information Technology Services." IDC MarketScape: Worldwide AI Services for State and Local Government 2025 Vendor Assessment, Oct. 2025. Accessed 12 Dec. 2025. [11]

"FedRAMP High Authorization: How It Helps." 1Kosmos. Accessed 12 Dec. 2025. [12, 37, 38]

"Generative AI Use and Management at Federal Agencies." U.S. Government Accountability Office, July 2025. Accessed 12 Dec. 2025. [14, 15, 16, 17, 39]

"Securing Artificial Intelligence." Booz Allen Hamilton. Accessed 12 Dec. 2025. [18, 44, 45, 46]

"NIST AI Risk Management Framework." National Institute of Standards and Technology, 2023. Accessed 12 Dec. 2025. [20, 34]

"NIST Says Risk Management Is Central to Generative AI Adoption." GovCIO Media, 19 Nov. 2025. Accessed 12 Dec. 2025. [21]

"Agency AI Use Doubled in 2024, GAO Finds." Nextgov/FCW, 29 July 2025. Accessed 12 Dec. 2025. [22, 26, 27, 28, 29]

"State Department Unveils AI Strategy to Modernize Diplomacy." MeriTalk, 2 Oct. 2025. Accessed 12 Dec. 2025. [23, 43]

"CrowdStrike Charlotte AI Achieves FedRAMP High Authorization." CrowdStrike, 25 Nov. 2025. Accessed 12 Dec. 2025. [30]

"NIST AI Risk Management Framework Playbook." National Institute of Standards and Technology, 2024. Accessed 12 Dec. 2025. [31, 32, 33]

"AI Use Case Inventory." Department of Homeland Security, 2024. Accessed 12 Dec. 2025. [47, 48]

Additional Reading

"C3 AI Achieves FedRAMP Authorization." C3 AI, 11 Dec. 2025. Accessed 12 Dec. 2025.

"Department of State AI Inventory 2024." U.S. Department of State, 2024, 2021-2025. Accessed 12 Dec. 2025.

"Impact of AI on NOCs." Jupiter Systems. Accessed 12 Dec. 2025.

"IT Outsourcing Trends for 2025." Transition Technologies MS. Accessed 12 Dec. 2025.

"Why Federal AI Needs a Secure Path to Scale." Leidos. Accessed 12 Dec. 2025.

About the Authors

Gregory A. Garrett is the Chief Operating Officer and Chief Innovation Officer for Evolver. He leads all of Evolver's technology thought leadership and technology business units, which provide provide advanced IT, AI, cybersecurity, electronic security, and eDiscovery services to federal agencies and Fortune 500 companies.

With over 25 years of executive experience and a background as a CIO, CTO, CISO, and COO, he has managed more than $40 billion in global tech contracts and served in key roles across both large and mid-sized firms.

A retired U.S. Air Force Colonel and accomplished author, Garrett has also contributed extensively to industry thought leadership through 24 books, more than 150 articles, expert testimony, and over 250 speaking engagements.

Rahul Johri is Evolver's Vice President of Digital Transformation and leader of the Evolver Innovation Center, bringing 30 years of experience delivering complex IT programs across the federal government. He leads enterprise modernization initiatives that integrate artificial intelligence, cloud, SecDevOps, and biometrics to drive mission impact at scale.

Rahul's background spans executive leadership in business operations, solutions architecture, and large-scale program delivery. Prior to Evolver, he served as Senior Director for National Security at ASRC Federal, with DHS portfolio P&L responsibility and leadership of major capture and solution efforts supporting TSA, FEMA, U.S. Coast Guard, USCIS, and the U.S. Secret Service.

Dr. Brian McElyea is an accomplished cybersecurity executive recognized for leading largescale digital transformation initiatives and building high-performing Cyber Centers of Excellence that drive mission-aligned security outcomes.

With over two decades of leadership across healthcare, defense, and federal sectors, he has guided global organizations through the complexities of risk management, compliance, and modernization with measurable results.