

Agentic AI in Government: Moving from Experimentation to Operational Impact

CONTENT PRODUCED BY



HOSTED BY

Deloitte.

This white paper summarizes an off-the-record roundtable discussion among government and industry leaders regarding the future of agentic AI in the public sector.



Introduction

Federal agencies are entering a new phase in their artificial intelligence journey. Over the past several years, government organizations have explored machine learning, robotic process automation, predictive analytics, and more recently, generative artificial intelligence. Initial efforts focused largely on experimentation, proof-of-concept activities, and isolated pilots designed to test the capabilities of rapidly evolving technologies. Today, however, federal leaders are increasingly focused on a different challenge. The question is no longer whether AI can provide value, but rather how agencies can operate AI capabilities in ways that improve mission performance, reduce administrative burden, maintain public trust, and deliver measurable outcomes at scale.

This roundtable, brought together federal technology leaders, chief data officers, cybersecurity professionals, operational executives, mission owners, and industry leaders to discuss the future of Agentic AI across government. Participants explored opportunities for near-term value creation, the maturity of emerging technologies, governance considerations, workforce implications, funding challenges, infrastructure requirements, and the practical realities of transitioning AI initiatives from pilot projects to enterprise capabilities.

While perspectives varied based on agency missions, several themes emerged throughout the discussion. Participants agreed that AI is already creating meaningful operational value in carefully selected use cases. They also agreed that many of the obstacles to adoption are no longer technical. Instead, organizations are confronting challenges involving governance, accountability, data readiness, workforce adaptation, procurement, budgeting, and organizational change. The conversation revealed a federal community increasingly focused on implementation rather than experimentation.

Where Agentic AI Is Creating Meaningful Value Today

The discussion began with a focus on practical applications and immediate opportunities. Participants described a wide range of use cases where Agentic AI is already demonstrating value or is expected to create meaningful impact in the near term.

One recurring theme involved operational decision support. Agencies are increasingly seeking ways to aggregate information from multiple enterprise systems and deliver relevant insights directly to supervisors, managers, and operational personnel. Rather than requiring employees to navigate dozens of dashboards and data sources, AI-enabled systems can synthesize information into concise summaries, identify anomalies, recommend potential courses of action, and provide context to support human decision-making, while responsibility for operational decisions remains with agency personnel.



Participants described environments where supervisors oversee thousands of personnel, manage complex logistics operations, or coordinate activities across geographically distributed organizations. In these scenarios, AI can serve as a force multiplier by helping personnel process information more efficiently and focus attention on emerging issues requiring human judgment.

Customer service and help desk environments also emerged as promising opportunities. Federal organizations often operate numerous service centers that interact with government employees and constituents. Participants discussed the potential for AI-enabled assistants to guide users through common administrative processes, automate requests, complete forms, answer routine questions, and direct individuals to the appropriate resources. These use cases offer measurable outcomes while presenting relatively low operational risk.

Knowledge management represented another area where participants noted they were seeing substantial benefits. Government organizations possess repositories of historical documents, reports, investigations, policies, and operational records. Much of this information remains difficult to access or analyze efficiently. Participants described efforts to use AI to enrich metadata, identify relationships among documents, summarize historical records, and make institutional knowledge more discoverable across the enterprise.

Several participants also highlighted opportunities within software development and engineering environments. AI-assisted coding, documentation generation, testing, and code analysis are helping organizations accelerate development activities while improving consistency and reducing administrative burdens on technical teams.

Across these examples, participants emphasized a common principle. Effective use cases are not replacing human decision-makers. Instead, they are augmenting human capabilities and enabling employees to operate more efficiently.

Is Agentic AI Overhyped or Ready for Government?

As enthusiasm surrounding AI continues to grow, participants engaged in a candid discussion regarding the maturity of current technologies. The consensus view was nuanced. Many participants rejected the notion that AI is merely hype, but they also cautioned against unrealistic expectations.

Many agencies have already demonstrated that AI can create operational value. However, participants noted that public conversations frequently emphasize theoretical capabilities rather than implementation realities. Demonstrating a capability in a controlled environment is fundamentally different from deploying that capability within a secure, governed, mission-critical federal environment.

Several participants observed that current discussions often focus excessively on models while overlooking the broader ecosystem required for effective implementation. Data integration, workflow design, governance, security controls, monitoring, workforce readiness, and operational accountability all play critical roles in determining whether AI initiatives are effective.

Participants generally agreed that the technology is mature enough to support many operational use cases today. At the same time, they cautioned that agencies should be skeptical of claims suggesting fully autonomous systems can reliably perform high-consequence government functions without meaningful human oversight.

Rather than pursuing ambitious visions of complete automation, participants advocated for targeted deployments that solve specific operational challenges and generate measurable outcomes. This approach allows agencies to build confidence, develop proficiency, and establish governance frameworks before expanding into more complex use cases.

Understanding ROI and Measuring Achievement

One of the significant themes of the roundtable involved measuring value. As agencies move beyond pilot projects, leaders are increasingly being asked to justify investments, demonstrate results, and establish sustainable funding models.

Participants acknowledged that traditional approaches to measuring return on investment may not fully capture the benefits and costs associated with AI deployments. While time savings remain important, organizations should also account for implementation costs, infrastructure expenses, token consumption, governance requirements, security controls, training activities, and ongoing operational support.



Several participants stressed the importance of establishing baseline measurements before introducing AI-enabled capabilities. Without understanding current performance levels, agencies cannot accurately quantify improvements or demonstrate meaningful outcomes.

The discussion also highlighted emerging concerns regarding AI economics. Current pricing models remain fluid, and many participants expressed concern that future costs may increase significantly as providers seek sustainable business models. Leaders noted that agencies should begin planning for environments in which AI consumption becomes a recurring operational expense rather than a temporary innovation initiative.

At the same time, participants recognized that some benefits may be difficult to quantify through purely financial measures. Improved service delivery, faster decision-making, enhanced employee productivity, reduced administrative burden, and increased organizational agility all stand to contribute substantial value even when they do not immediately translate into budget reductions.

Participants repeatedly emphasized that future effectiveness may depend on the government's ability to establish more sophisticated approaches to measuring AI value and aligning investments with mission outcomes.

Data Readiness as the Foundation for AI Effectiveness

No topic received more consistent attention than data. Participants repeatedly observed that effective AI implementation depends far more on data quality and accessibility than on model selection.

Many agencies continue to operate within environments characterized by fragmented systems, inconsistent standards, legacy applications, and organizational silos. These conditions create significant challenges for AI adoption because intelligent systems depend upon access to reliable, well-governed information.

Participants described numerous challenges involving data ownership, privacy restrictions, access controls, and interoperability. Investigative organizations noted that relevant information frequently exists across multiple systems governed by different legal and operational requirements. Human resources, logistics, financial management, and operational systems often maintain separate repositories that were not designed to work together.

The discussion reinforced a growing understanding across government that AI readiness begins with data readiness. Organizations that invest in data governance, metadata management, interoperability, and information quality are significantly better positioned to achieve meaningful results.

Participants also highlighted the importance of understanding data rights and usage permissions. Agencies should clearly define what information can be accessed, how it can be used, and what safeguards must be applied before integrating data into AI workflows.

Without strong data foundations, participants warned that even the most advanced models will produce unreliable outcomes and undermine trust in AI-enabled systems.

Governance, Accountability, and Trust

As AI systems become more deeply integrated into government operations, governance is emerging as one of the leading areas of focus.

Participants described governance as operating across multiple layers. Builders and developers require guidance regarding approved tools, acceptable use, security requirements, and development practices. Platform owners should monitor system performance, oversee integrations, and manage operational risk. Mission owners should establish accountability for outcomes, determine appropriate use cases, and confirm alignment with organizational objectives.

The conversation emphasized that governance should not be treated as an afterthought. It should be embedded throughout the lifecycle of AI systems, from design and development through deployment, monitoring, and retirement.

Several participants noted that some agencies are effectively building governance models while simultaneously deploying new capabilities. This can create tension between innovation and oversight. Organizations should move quickly to capture value while recognizing that governance controls will need to evolve alongside AI implementation. Rather than establishing static policies, agencies should continuously assess and refine governance practices as technologies, risks, and mission requirements change.



Participants also discussed the importance of clearly defining vendor and organizational responsibilities as AI technologies evolve. This included the need for greater transparency into AI software components, such as software bills of materials (SBOMs), as well as clearly established processes for addressing vulnerabilities, applying patches, and maintaining systems over time. Several participants noted that organizations should not bear sole responsibility for rapidly responding to emerging risks, particularly cybersecurity vulnerabilities, and emphasized that vendors have an important role in maintaining and securing the technologies they provide.

Trust emerged as a central consideration. Employees, leadership, and stakeholders must understand how AI systems operate, where outputs originate, and how decisions are being influenced. Participants stressed that transparency and explainability are essential for maintaining confidence in AI-enabled operations.

Ultimately, participants concluded that governance frameworks will play a decisive role in determining whether AI initiatives can effectively scale across government.

Moving from Pilot Projects to Enterprise Capabilities

The transition from pilot projects to operational deployment generated extensive discussion. Participants agreed that many AI initiatives struggle not because of technological limitations but because organizations underestimate the complexity of scaling effective pilots.

Pilot environments typically operate under controlled conditions with limited users, defined data sources, and manageable risk profiles. Enterprise deployment introduces significantly greater complexity. Agencies must address security reviews, compliance requirements, operational support, workforce training, monitoring capabilities, and integration with existing systems.

Participants also highlighted challenges involving organizational ownership. Effective scaling requires clear accountability for both technology platforms and mission outcomes. Without defined ownership structures, promising initiatives frequently stall before reaching operational maturity.

The conversation revealed growing recognition that AI deployment should be viewed as an ongoing operational capability rather than a technology project with a fixed endpoint. Organizations should establish processes for continuous improvement, performance monitoring, governance updates, and adaptation as technologies evolve.

Agencies that effectively navigate these challenges will likely be better positioned to transform isolated achievements into sustainable enterprise capabilities.

Workforce Transformation and Human-AI Collaboration

The workforce implications of Agentic AI remain among the most important and misunderstood aspects of federal AI adoption.

Participants consistently rejected narratives suggesting that AI will replace large portions of the federal workforce. Instead, they described a future characterized by collaboration between humans and intelligent systems.

Many of the leading use cases involve reducing administrative burdens and enabling employees to focus on higher-value activities. AI can assist with information gathering, document analysis, workflow management, data enrichment, and routine decision support while leaving critical judgment and accountability in human hands.

This transition requires significant workforce preparation. Employees should understand how AI systems operate, recognize their strengths and limitations, validate outputs, and apply appropriate oversight. Agencies will likely need new training programs, revised operating procedures, and updated workforce development strategies.

Participants emphasized that organizational culture may be as important as technical skills. Employees should be enabled to understand that AI is intended to augment their capabilities rather than replace them. Building trust and confidence is essential for driving adoption and achieving sustainable results.

The Emerging Federal Operating Model

As the discussion concluded, participants reflected on the future of government operations in an AI-enabled environment.

A common vision emerged. AI will likely become increasingly embedded within everyday workflows, supporting information discovery, decision-making, coordination, and operational execution. Human knowhow will remain essential, but employees will increasingly rely on intelligent systems to process information, surface insights, and manage routine activities.

An AI enabled federal operating model will require stronger data governance, more sophisticated performance measurement, integrated security frameworks, and greater collaboration across organizational boundaries. Agencies are encouraged to establish funding models that support innovation while maintaining accountability. They will likely also need governance structures capable of evolving alongside rapidly changing technologies.

Importantly, participants emphasized that effective implementation depends on leadership. Technology alone will not drive transformation. Leaders should establish priorities, manage expectations, build trust, and align AI initiatives with mission objectives.



Conclusion

Federal agencies are entering a pivotal period in the evolution of artificial intelligence. The conversation demonstrated that government has moved beyond curiosity and experimentation. Agencies are now confronting the practical realities of implementing and governing AI systems, demonstrating measurable return on investment, preparing the workforce, and integrating AI into mission operations.

The discussion revealed widespread optimism regarding the potential of Agentic AI to improve government operations. At the same time, participants recognized that realizing this potential requires disciplined execution. Success depends on strong data and technology foundations, adaptive governance frameworks, human oversight, workforce readiness, security, and sustainable funding models that support implementation over time.

Participants also emphasized that workforce transformation extends beyond increasing AI adoption. As agencies incorporate AI-generated code and automated development tools into software engineering workflows, they will continue to need technical professionals with the expertise to review, validate, maintain, and improve AI-generated outputs. Building these capabilities will be essential to safely and effectively operationalizing AI across government.

Finally, participants noted that AI technologies will continue to evolve rapidly, requiring agencies to plan for ongoing change rather than one-time implementation. Governance models, technical architectures, and operational processes should be designed to accommodate new models, evolving capabilities, and changing security requirements while minimizing disruption to mission operations.

The agencies that may achieve the greatest results will not necessarily be those deploying the most advanced models. They will likely be the organizations that establish strong data, governance, and operational foundations, maintain human oversight of mission-critical decisions, continuously adapt as the technology evolves, and remain focused on delivering measurable mission outcomes.

The future of government AI will be defined not by experimentation alone, but by the ability to translate emerging capabilities into sustainable, secure, and measurable mission impact.