Microsoft

The federal AI playbook: Mission transformation for the AI era

Artificial Intelligence (AI) will likely be the most impactful technology of this era. AI's effect on essential workflows, from content creation to search and analysis to decision support, is already being felt across the government.

Government initiatives—from the 2020 AI in Government Act to the Biden Administration's recent AI Executive Order—speak to the fundamental tension that agencies must navigate as AI becomes more available and capable. Innovation vs. risk is always a difficult balance, but AI's potential is only starting to be understood.

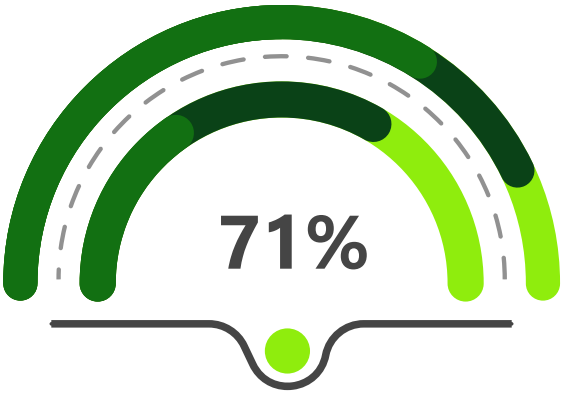

Government decision makers who believe the benefits of generative AI for agency operations outweigh the risks[1]

The path forward must include guardrails to ensure that AI is applied in an ethical, secure, transparent, and human-centric manner across all sectors. At the same time, it's crucial to enable the rapid innovation needed for economic prosperity and national security.

As government agencies continue experimenting with AI use cases, one important question stands out: how will constituents benefit from the technology? From climate research to filing taxes to veterans' healthcare, AI can be an absolute game-changer in how agencies operate and deliver critical services to the people that need them. Along with understanding how AI can transform operations, it's crucial to consider the impact on people, processes, and systems.

1 "Government gears up to embrace generative AI," October 17, 2023. FedScoop.

How AI accelerates mission outcomes

One capability that holds incredible public sector promise is AI's ability to analyze enormous amounts of data at unheard-of speeds and make predictive recommendations based on the insights. Generative AI––a form of AI that can produce various types of content, including text, imagery, audio, code, and synthetic data––has the potential to transform mission work, in areas such as:[1]

Document creation and review: Imagine handing over a lengthy regulatory document and simply instructing the AI, "Identify the essential regulatory requirements and provide a concise summary." With some thoughtful adjustment of these instructions, also called "prompt engineering," a Generative AI system can deliver user-friendly output in seconds, allowing people to focus on higher-level decision making. Human reviewers, as always, have the last word, ensuring the output is appropriate and accurate.

Customer experience: The fusion of customer experience (CX) and AI can transform the way federal agencies interact with the public. An AI-powered chatbot can pull information from multiple systems of record to understand the user's prior interactions with the agency and even recommend next steps. A natural language interface makes accessing government services easier, faster, and less intimidating. This enables agencies to deliver personalized services that meet the needs and expectations of every user.

Research and intelligence analysis: From correlating clinical trial results to identifying connections within open source intelligence, Generative AI can help researchers, scientists and analysts make the best use of available information. It can identify supporting and contradictory evidence for a working hypothesis, summarize evidence and arguments from multiple sources, suggest alternatives, and highlight key information, all saving significant time and effort.

Data harmonization: Creating a unified data set from multiple studies across a myriad of siloed systems requires cleaning, de-duplicating, and melding it into a common format. This takes a massive amount of time and processing power. Using Generative AI to identify and link specific data can make the process exponentially faster, freeing users to focus on solving mission-related challenges.

These examples are just scratching the surface of what might be possible with GenAI. There are potential use cases within use cases, which is why secure experimentation is so critical as agencies assess AI's future value.

Building the AI-enabled government workforce

Many of the use cases above drive mission speed, efficiency and accuracy gains by doing one simple thing: helping agency personnel do their jobs better.

IT leaders who believe AI will:[1]

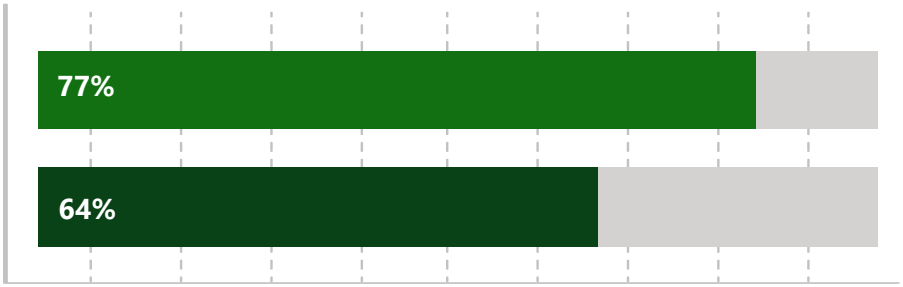

This point is critical: AI isn't replacing humans. It's helping people create, innovate, and achieve more than humans or technology on their own could accomplish. Generative AI can augment people's capabilities to speed up the processing of tax returns for the Internal Revenue Service (IRS) or help the Department of Veterans Affairs decrease the time it takes to process claims.

What many don't realize is that many of these AI-powered capabilities are already available and are starting to be embedded within productivity applications that federal personnel use every day. Imagine working on various tasks—from building a slide deck to analyzing financial figures—and having AI-powered features on-demand. That day is here.

To make the most of these AI-powered copilots, federal workers need not only to understand how to use these tools, but also comprehend their strengths and weaknesses. Users may need to ask the question in a different way to get the best result. Training is vital to help users learn how to refine their prompts. It's worth the effort, as integrating AI solutions seamlessly into everyday tasks and processes can drive innovation and productivity across federal agencies.

Constructing an AI-ready infrastructure

Along with preparing people to make the best use of AI, the environment needs to be optimized to deliver performance with security. Cloud solutions, from native applications to storage and processing capabilities, need to scale to support AI's computing and data management requirements. An effective strategy starts with identifying business needs: instead of "technology for its own sake," it's crucial to spell out use cases for Generative AI. But it's just as important to consider future uses, which could need more resources. Taking the long-term into account allows more room for agencies to experiment with innovative uses for AI-powered solutions, which could result in huge benefits for the mission.

1 "Government gears up to embrace generative AI," October 17, 2023. FedScoop.

Many agencies already have a cloud platform in use, but for new AI workloads, it's important to consider which cloud is best for both the near-term use cases and future expandability. Look at platform providers in terms of support for AI-powered solutions, including their investment in supporting emerging technologies. Also, look at the multicloud management capabilities they provide, to enable a single point of control and visibility. And of course, take time to understand how they implement security at every level needed.

Scale cyber readiness with AI

AI advances have important implications for our nation's cyber posture and national security. Like any emerging technology, the rising capabilities of AI are available to friends and foes alike.

AI can play a crucial role in powering a Zero Trust framework. Real-time monitoring delivered through AI can quickly identify anomalies and potential threats, perpetually learning and adapting to emerging patterns. Similarly, AI can help detect unusual user behavior that might indicate malicious intent or compromised accounts, mitigating the risk of insider attacks.

Beyond faster, more accurate detection, AI-driven security automation also empowers agencies to streamline incident response, instantly taking pre-defined steps to protect vital systems and data.

AI can also help ensure that security policies are enforced consistently across the organization–– especially important in complex environments with a mix of remote and on-site workers. Overall, this kind of intelligent automation can significantly enhance an agency's cybersecurity posture and protect sensitive data, critical infrastructure, and citizen information from a wide range of threats.

Reservations regarding the security of AI tools are often rooted in misconceptions about the technology. Enterprise-grade AI represents a substantial departure from the generative AI tools commonly found in the public domain, such as ChatGPT.

Securing the enterprise

Unlike its consumer counterparts, enterprise AI is fortified with an array of security features to protect sensitive data and uphold compliance with government regulations. Additionally, enterprise AI requires transparency about how data is used: where it's stored, how it's processed, and what is logged—and what isn't.

The public sector must use systems that can be trusted to keep data safe in its own environment. Microsoft, for example, recommends setting up a private, secure Generative AI sandbox in the cloud. This allows users to experiment with—and securely chat with—the technology. This approach also gives organizations an isolated test environment to learn what AI can do while leveraging cloud security controls to maintain privacy

for model prompts, generated output, and source data sets. Privacy includes the model developer or host not reusing the data in any way, including to improve the general-purpose model. In plain language, "Your data is your data."

The foundational standards for measuring the risk of other emerging technologies can apply to AI as well. For example, the FedRAMP standards and the NIST RMF contain components that can serve as a fantastic foundation for ensuring secure and responsible AI deployment. The NIST AI Risk Management Framework (AI RMF) is intended for voluntary use and to improve the ability to incorporate trustworthiness considerations into the design, development, use and evaluation of AI products, services and systems.

These enhanced safeguards provide assurance that integrating AI into mission workflows does not compromise data integrity. At the same time, greater confidence in the security of the system empowers an environment ripe for heightened productivity and innovation.

Ensuring AI trust through transparency

AI risk goes beyond security to something more fundamental: trust through transparency.

For AI to have any value, government personnel have to be able to trust the accuracy, reliability, and ethics of its conclusions. User experiences should also be designed to mitigate AI risks and encourage the responsible use of the output without over-relying on the technology. The common thread in addressing these concerns is transparency.

AI-based systems––even ones using large language models whose internal processing can be difficult to understand––should show how they derived an answer. It's critical for users to confirm that the sources used are, in fact, reliable, and that the conclusions drawn by the system make logical sense. In other words, AI systems should show "answers plus evidence," and users should be trained to understand that evidence.

Both the creators of AI technologies and users have the responsibility to ensure its ethical use. From a development standpoint, for example, Microsoft is careful to ensure that AI ethics are built-in to both the design and practices of these tools and that other organizations creating derivative products on top of this technology do the same.

Governance ensures confidence

Enterprise AI should also allay concerns about proprietary data being used to train models. This should only happen when specifically called for; otherwise, only non-proprietary data is used for training. How data is used is a matter of governance and policy.

To tackle all of the challenges of the ethical use of AI, it's critical for government agencies to consider bias when collecting data and training models. However, all bias will likely never be eliminated from all models. Because of this, it is equally vital to put tools in place that filter content within models to ensure transparency and compliance. These tools enable users to remove sensitive terms in responses. Again, the agency's governance determines how these filters are applied, giving humans full control of the system's output.

As the federal government continues to invest in and use AI tools, decision-makers in the federal government must have the appropriate training to ensure this technology is used responsibly and ethically. With AI training, federal agency leaders will have the expertise needed to ensure this technology benefits the public and mitigates potential harms, such as bias and discrimination.

Starting Steps for Success

As the future of work moves toward an AI-powered digital workspace, it's becoming increasingly critical for government agencies to embrace this change to stay ahead of the curve and seize opportunities to enhance efficiency, drive innovation, and improve citizen services. As governments look to integrate AI into their daily operations, the following steps should be guiding principles to ensure success:

1. **Identify low-risk use cases** to experiment with AI tools and embrace the technology at a comfortable pace;
2. **Build your own secure experimentation sandbox** where AI can be internally tested, evaluated and refined before public deployment;
3. **Leverage AI that's embedded within existing enterprise** tools to ensure high-grade security;
4. **Lean on existing risk standards** to ensure AI compliance – without having to reinvent the wheel.

Agencies can potentially pave the way for capturing the full value of AI by identifying where the biggest opportunities lie and supporting the adoption of AI technologies in an ethical and secure manner.

AI innovation has the potential to transform all areas of government. This unlimited future is why Microsoft now empowers agencies with Azure OpenAI Service in Azure Government. With the accelerating pace of AI adoption worldwide, it's essential for agency leaders to consider how, not when, to implement AI effectively, safely, and securely.

Microsoft

To learn more about bringing the power of Generative AI to government agencies, visit Microsoft's Federal AI Hub.