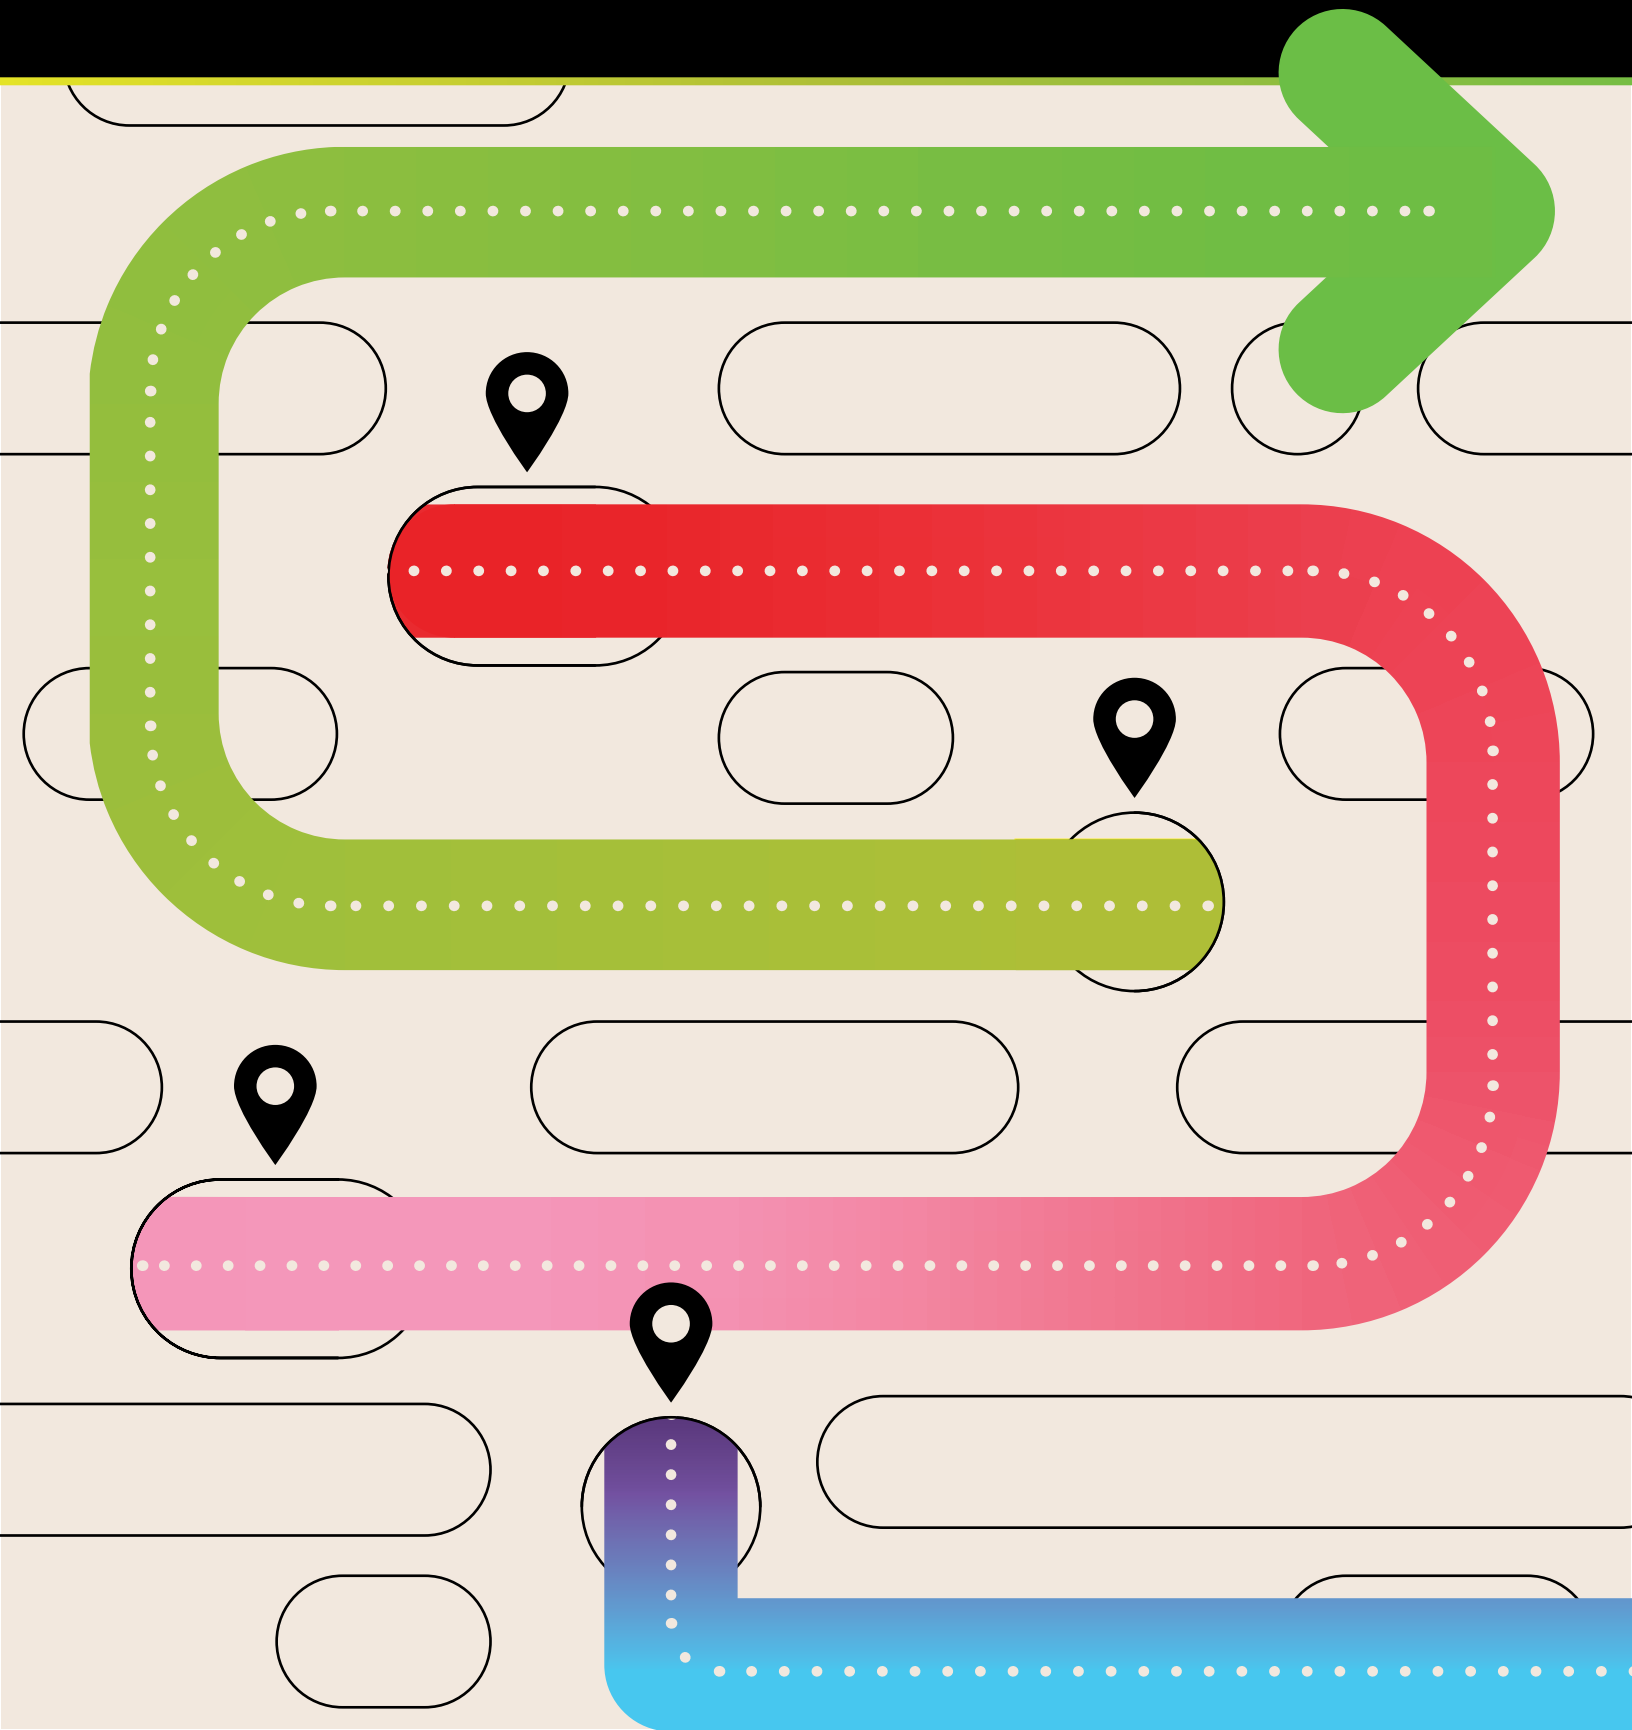
AI: Where We Are, Where We're Going



Welcome to the World of Responsible AI

Artificial intelligence (AI) is finding its way into more and more applications, offering exciting new ways to analyze data, respond to the public and streamline internal operations. In government, the past year has seen widespread adoption and ongoing refinement of the technology. As agencies gain experience, they're learning both what AI does well and where it needs to improve.

That's led to an emphasis on "responsible AI," i.e., the steps governments should take to keep their AI trustworthy, unbiased and secure. In this third and final installment of our 2024 AI guide series, we look at how governments are addressing the need to make AI safe and productive.

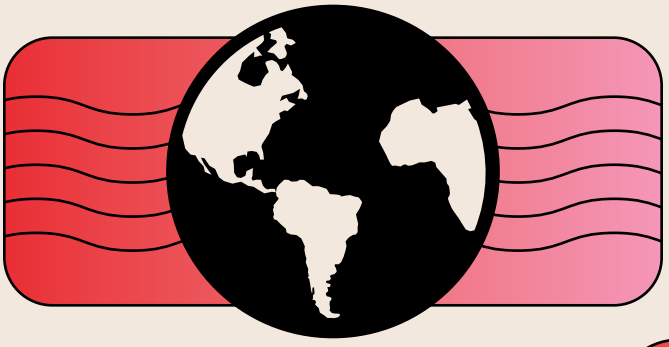


Table of Contents

3 AI's Evolving Agenda

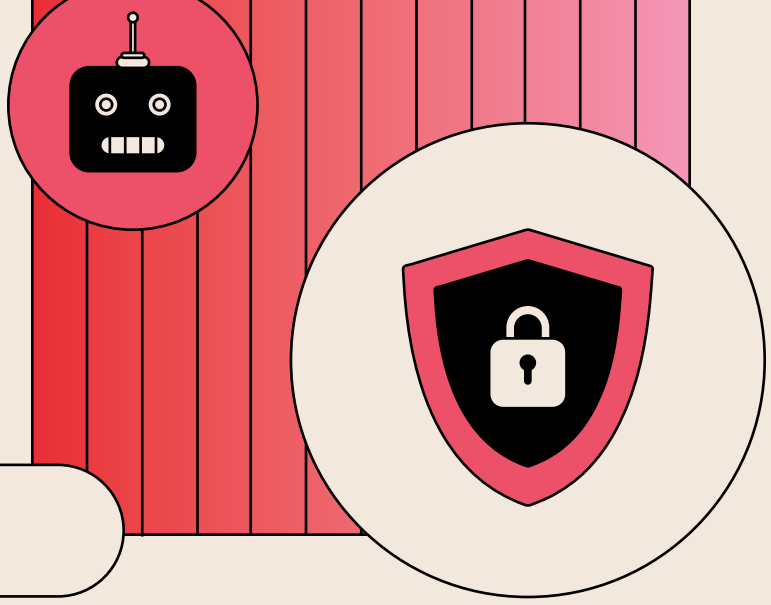
6 Transforming Government With AI

7 On the Road to Responsible AI

9 Skills Benchmarking Lays the Groundwork for AI Success

10 What's Ahead

An interview with Alan Fuller, CIO of Utah



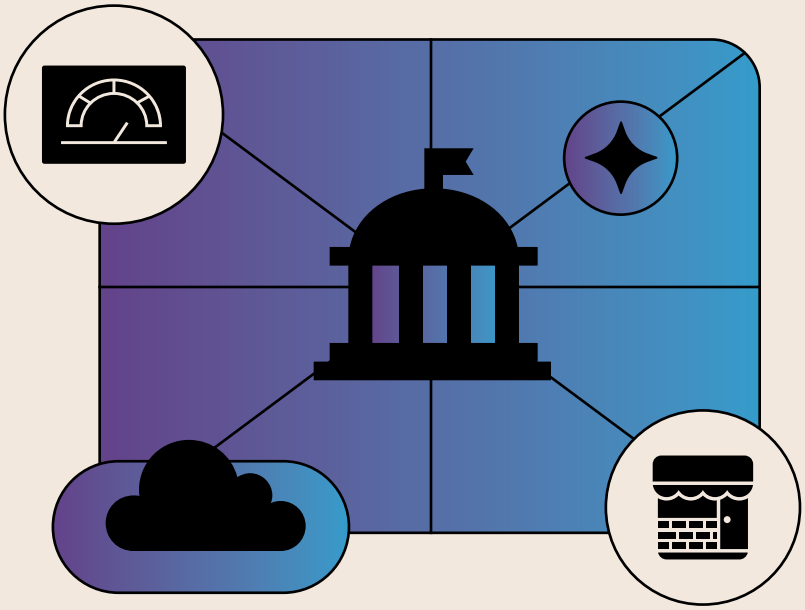
AI's Evolving Agenda

Understanding the safe and responsible use of AI remains one of the biggest priorities for agencies heading into 2025 (see article, p. 7), but it is not the only one. Here are some other issues high on agency agendas.

1. Strengthening the AI Marketplace

One of the challenges with any new technology is understanding how to procure it in a way that both supports the emerging marketplace and protects the government's interests. Here are some of the questions that the U.S. Office of Management and Budget (OMB) raised in a public notice earlier this year:

- How may standard practices and strategies of federal procurement best be used to reflect emerging practices in AI procurement?
- How can OMB promote robust competition, attract new entrants into the federal marketplace (including small businesses) and avoid vendor lock-in?
- Should the federal government standardize how it assesses the benefits and trade-offs among in-house AI development, contracted AI development, licensing of AI-enabled software and use of AI-enabled services?
- How might government develop and communicate metrics to enable performance-based procurement of AI?



2. Improving the Constituent Experience

Agencies continue to explore AI's potential to address another high priority: making government services more effective, efficient and personalized. For example, on July 31, 2024, the U.S. General Services Administration hosted the Federal AI Hackathon, in which 250 people used AI and cloud tools to optimize government services. Here are the winning projects:

- **First place:** an original interface and AI-powered search engine for the National Science Foundation's awards search website.
- **Second place and People's Choice:** an AI-powered assistant that returned easy-to-understand article summaries from AskUSDA, the U.S. Department of Agriculture's content hub.
- **Third place:** an AI-powered chat interface so users could ask questions and get information from the Census Bureau's data platform.

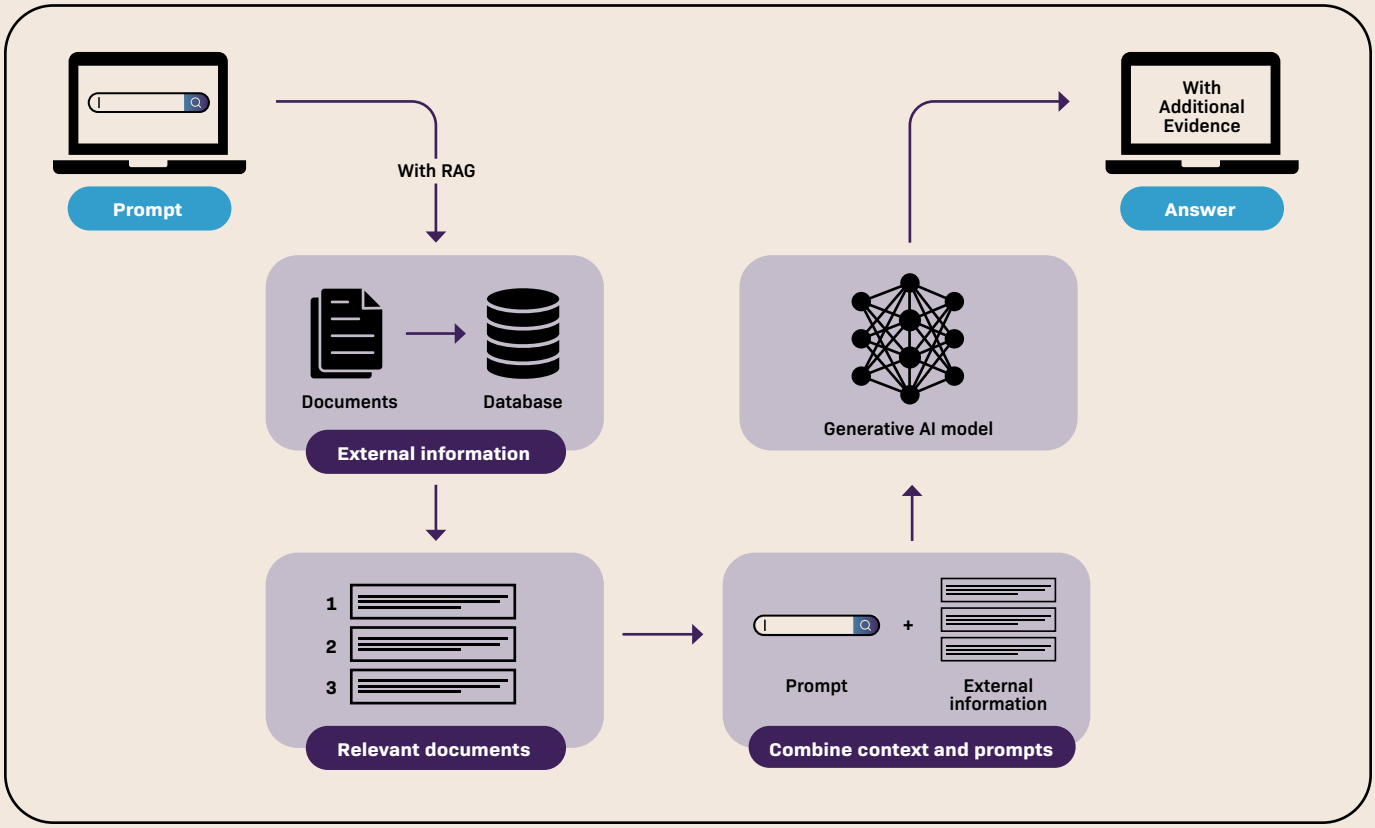
And here are criteria the judges used:

- **Innovation:** How novel and creative is the solution?
- **Functionality:** Does the solution work seamlessly with AI tools?
- **Usability:** Is the website user-friendly for both humans and AI?
- **Scalability:** Is this a viable solution to scale to similarly situated websites?
- **Impact:** How significantly does the redesign improve the interaction with government services?

3. Boosting GenAI's IQ

For government users, a common weakness with commercial generative AI (GenAI) tools is that they were trained using general datasets that don't necessarily align with government use cases. One technique to address that concern is retrieval-augmented generation (RAG), according to a report from the U.S. Government Accountability Office. RAG supplements a commercial large language model (LLM) with contextual information from an agency's databases.

Here is an example of how RAG works:



4. Busting AI Myths

Here are some of the most common myths around AI among local governments, according to a recent report from the National Association of Counties (NACo):

- ✕ *My county can avoid the use of GenAI by never allowing employees to use the technology.*
- ✕ *AI is a new technology that society has not previously used.*
- ✕ *GenAI models are always accurate and do not need their outputs verified.*
- ✕ *I can turn off GenAI features in applications if I don't want to use them.*
- ✕ *All GenAI tools are designed using the same underlying open source model.*
- ✕ *GenAI tools will return the same answer every time, so long as the prompt you offer is written the same exact way.*
- ✕ *AI and GenAI can steal any data that I enter into the model.*

5. Building an AI Curriculum

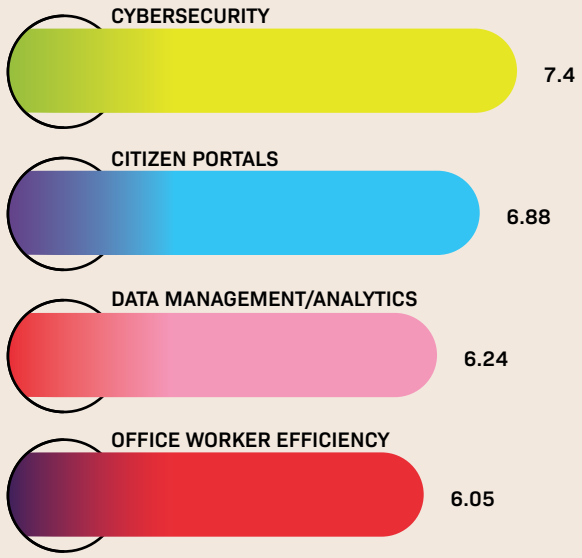
Agencies are scrambling to train employees to understand both the opportunities and the risks that come with AI. Here are some of the AI-related courses the IT Modernization Centers of Excellence offered to federal employees in 2024:

- **Buying AI: Government Contracts 101.** Basic federal procurement policies and requirements as they relate to AI.
- **Risk Management and Ethics.** Ethical considerations and risk mitigation measures essential to responsible AI acquisition.
- **Data Privacy Considerations.** General and AI procurement-specific data privacy concerns that are crucial to an effective and compliant procurement strategy.
- **AI Safety and Robustness.** Factors AI developers should evaluate when designing AI systems for safety, including how to address biased inputs and explainability issues.
- **AI Auditing.** How to audit AI systems so they adhere to performance, safety and ethical standards.

Cybersecurity Tops List of AI Use Cases

Here are the top four AI use cases related to IT operations, according to a recent survey of state technology directors:

HOW WOULD YOU RANK THE USE CASES YOUR STATE IS USING/FORESEES IN AI?

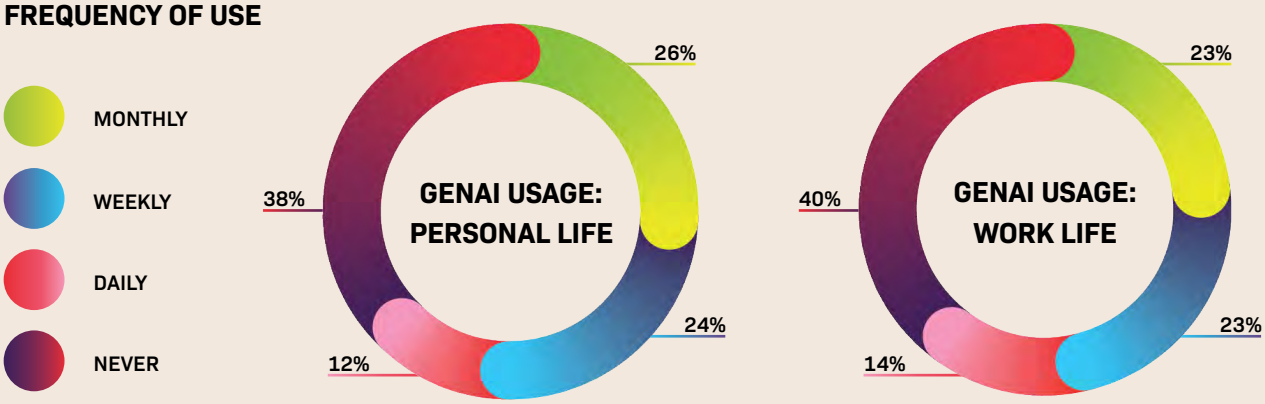


(SCALE OF 1-10: WEIGHTED AVERAGES)

GenAI: It's Personal

A recent survey by the NACo found that employees are using GenAI at home and at work in nearly equal measure:

FREQUENCY OF USE



Transforming Government With AI

An interview with Mark Hughes, Chief Strategy Officer; Stephen Moore, Chief Technology Officer; and Chris Steel, AI Practice Lead, AlphaSix



AI is changing the way government does business. Through automation, predictive analytics and advanced data analysis, AI is set to enhance operations and service delivery, said Chris Steel of AlphaSix, which provides data management platforms and tools for data analysis.

"Its transformative potential in government operations is vast, ranging from predictive analytics in defense and law enforcement to automating benefits processing, and everything in between," he said.

And it's time for agencies to get started, said Mark Hughes, also of AlphaSix.

"You don't have to start large if people are interested in playing with the tools and seeing what they can do," he said. "It easily scales as their needs scale."

Getting Practical

Agencies are using AI for a range of applications, Hughes said, including anticipating maintenance needs in public infrastructure, automating routine administrative tasks and using natural language processing for more responsive chatbots.

"Some of these innovations not only streamline processes but they make government services more accessible and responsive to citizen needs," Steel said.

AI also plays a significant role in cybersecurity. For instance, at a large federal agency, AlphaSix used machine learning (ML) and AI to create a complex anomaly-detection tool to protect sensitive customer data. "This not only secures critical data, but streamlines operations," Steel explained.

Overcoming Hurdles

Preparing to adopt AI takes work, AlphaSix's Stephen Moore said: "The first thing is just knowledge. You need [someone] who understands the technology behind the buzzwords." Because there's a shortage of skilled AI professionals, agencies may have to bring in vendors and consultants who can provide that expertise.

Another issue is having data that is AI-ready. "If your data's spread all over the place in a bunch of stovepipe systems, it's going to be hard to unlock the true value of looking at a higher level across datasets," Moore explained.

Connecting all those silos unleashes vast amounts of data, often on systems that aren't designed to handle those volumes, forcing analysts to summarize or use fewer datasets. "They're missing the input that they would need to get the output, the analysis they're looking for," Moore said.

AlphaSix brings 13 years of experience in developing data tools and platforms to help agencies spend less time on data management and more on analytics, Hughes said.

Its solution is a data fabric, a system that starts with bringing in billions of records from diverse applications and making them available for analysis. The next step is to run queries and build the required analytics. "And the third phase is to add in AI and ML so you can start building models that can detect more in real time," Moore said.

AI is ready for its role, Steel added: "It's time to embrace AI as a transformative tool for government."

AlphaSix | Hewlett Packard Enterprise

On the Road to Responsible AI

AI has come a long way in the past few years, and agencies have learned a lot about best practices. Where once they were issuing blanket bans of GenAI tools such as ChatGPT, they're now finding safe ways to adapt the technology to their needs and reaping the benefits.

The importance of using AI responsibly can't be understated. To quote Executive Order 14110 on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, "Responsible AI use has the potential to help solve urgent challenges while making our world more prosperous, productive, innovative, and secure. At the same time, irresponsible use could exacerbate societal harms such as fraud, discrimination, bias, and disinformation; displace and disempower workers; stifle competition; and pose risks to national security."

The good news is there's a growing body of guidance on how to meet these goals, starting with the EO and OMB's implementation memo, M-24-10. Even better, a consensus is emerging in both the public and private sectors on what responsible AI requires.

What Is Responsible AI?

We hear a lot about "trustworthy" and "ethical" AI. For example, the National Institute of Standards and Technology defines trustworthy AI as: "valid and reliable, safe, secure and resilient, accountable and transparent, explainable and interpretable, privacy enhanced, and fair with harmful bias managed."

But Cansu Canca, Director of Responsible AI Practice at the Northeastern University Institute for Experiential AI, says the term "responsible AI" is more accurate. "In a way, 'responsible AI' is a shorthand for responsible development and use of AI," Canca said. 'Trustworthy AI' seems to direct the attention to the end goal of creating trust in the user. [That] circumvents the hard work of integrating ethics into the development and deployment of AI systems."

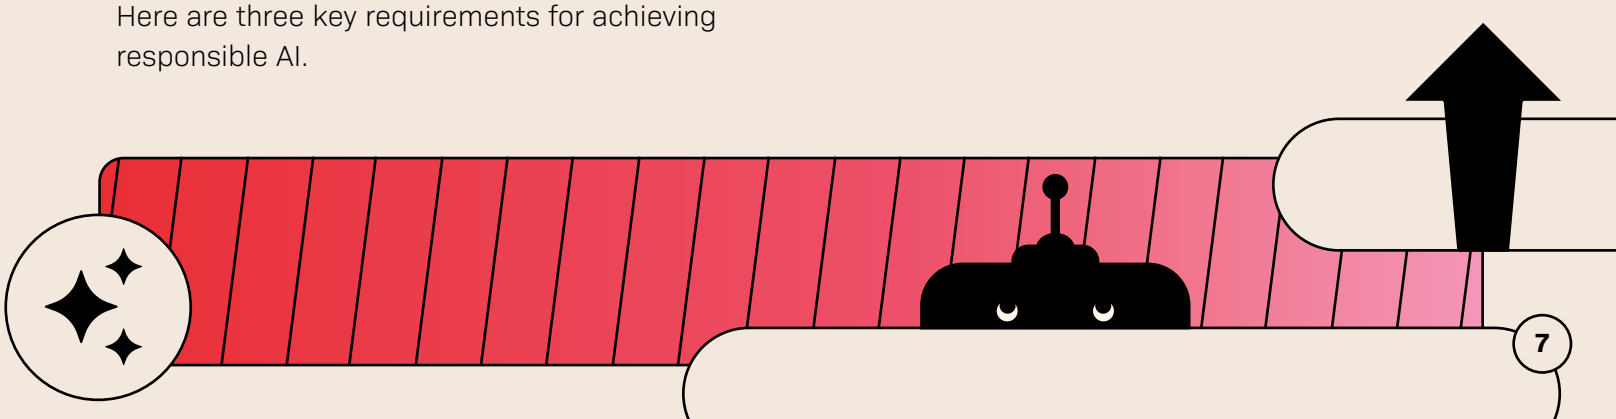
Here are three key requirements for achieving responsible AI.

1. Control the Data

People continue to revel in — and repeat — stories about GenAI's factual mistakes and absurd outputs. When an AI query yields advice to eat a rock every day, it raises skepticism about the whole technology.

Agencies have found that improving accuracy depends on limiting AI's frame of reference to accurate information relevant to task. This means moving away from publicly accessible LLMs and bringing the applications in-house, using "specialized" AI, trained on an agency's own data.

Vendors have expanded and simplified their offerings for this; it's not something agencies must develop themselves. But they will have to step up their data governance to make sure AI draws on information that's error-free, current, and appropriate to the task.



2. Dispel the Myth of the 'Black Box'

The terms “transparency and explainability” may have seemed mysterious a few years ago; now that AI is more familiar, they’re easier to understand. It turns out you don’t need to know how to write an algorithm to explain what data the AI accesses or what you’ve asked it to do. AI isn’t really a “black box” whose outputs are inscrutable.

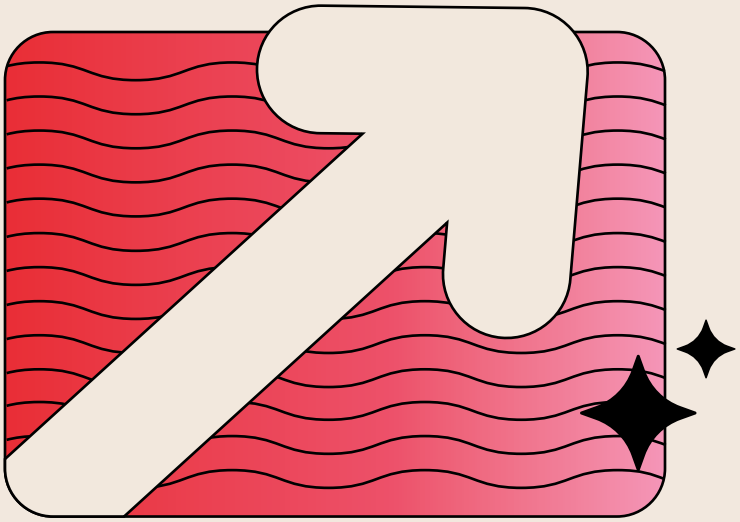
But agencies need to be proactive in explaining the limitations of AI-generated results, just as they would any other analytics. The most basic level of transparency is one that might be easy to overlook: *telling* users where AI is being used in products and services.

It also requires monitoring outputs to see where they may fall short, such as by reflecting bias in the data, and being up-front about the issues. The more you know about what goes into the AI, the more likely you will be to head off problems — and to mitigate them when they occur.

3. Don’t Forget Privacy and Security

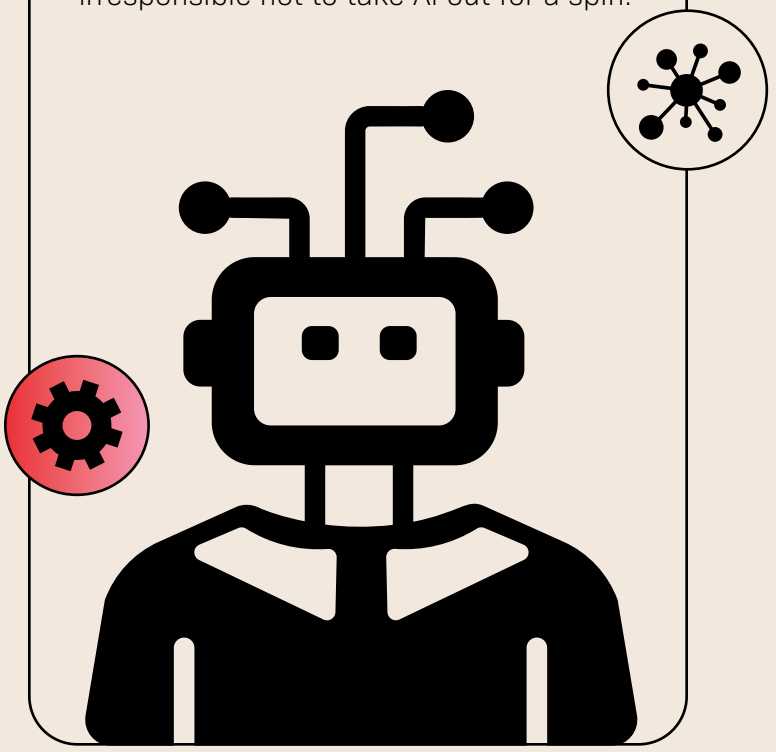
Government agencies collect a lot of data about the people they serve — much of it confidential or personally identifiable information (PII). As they apply AI to this data, they need to install guardrails that keep that information from leaking out. The technology’s ability to search and link data is useful in streamlining services, but an AI model trained on data that includes PII or other sensitive information can inadvertently reveal it in answer to a query.

Increasingly, AI tools can be trained to recognize PII and confidential information and block it from specific applications. But you can’t rely on AI to monitor itself. Before implementing any public-facing AI, agencies must be sure it doesn’t tell more than it should.



'Responsible' Is More Than Avoiding Risk

Although agencies often emphasize the “responsible” in “responsible use,” part of that responsibility is to *use* AI. Its ability to summarize, synthesize and organize vast amounts of data has demonstrated great potential for enhancing cybersecurity, chatbots, call centers and data analysis, among other uses. With these capabilities, it would be irresponsible not to take AI out for a spin.



Skills Benchmarking Lays the Groundwork for AI Success



"There are all kinds of executive orders out there telling us we need to benchmark our skills. Agencies need to know where their teams are, and then make a plan based on the data."

– Tony Holmes, Pluralsight

WATCH THE VIDEO

In the coming years, government agencies are looking to artificial intelligence (AI) and generative AI solutions to transform all aspects of their operations and services. The challenge is that they need to develop the workforce to support that vision.

The October 2023 Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence emphasized the importance of training employees in the benefits, risks and limitations of AI for their job functions. But workforce development involves more than just offering AI courses. It requires understanding the current skills gap and building a strategy for addressing it.

In this video interview, Tony Holmes, Practice Lead for Public Sector Solutions at Pluralsight, discusses how agencies can prepare their workforce for AI and other emerging technologies. Topics include:

- Using benchmarks to shape your agency's training programs
- Identifying employes who already have developed valuable GenAI skills
- Helping the workforce see learning as a part of their jobs

About Pluralsight

Pluralsight is the leading technology workforce development company that helps government agencies develop critical skills, improve processes and gain insights through data, and providing strategic skills consulting.

We help build technology skills at scale with expert-authored courses on today's most important technologies, including cybersecurity, cloud, AI, data science, and more. Our platform includes tools to align skill development with agency objectives, virtual instructor-led training, hands-on labs, skill assessments and one-of-a-kind analytics.

LEARN MORE

PLURALSIGHT

What's Ahead

An interview with Alan Fuller, Chief Information Officer, Utah Division of Technology Services



In March 2024, Utah's Gov. Spencer Cox signed a bill creating the nation's first Office of AI Policy (OAIP) to study advancements in AI technology, identify risks to consumers and issue policy recommendations governing AI. By June, OAIP was up and running, along with an AI learning lab where stakeholders can collaborate on new applications and the policy needed to regulate them.

Utah's agencies have long been at the forefront of using technology, including AI, to improve services and streamline operations. The state set a statewide GenAI policy at the end of 2023 and earned "A" grades from the Center for Digital Government's Digital States Survey in 2022 and 2024.

So what can agencies learn from Utah's experience? We spoke with Alan Fuller, Chief Information Officer of Utah's Division of Technology Services (DTS), about how the state is implementing AI, and what he sees ahead.

Generating Productivity

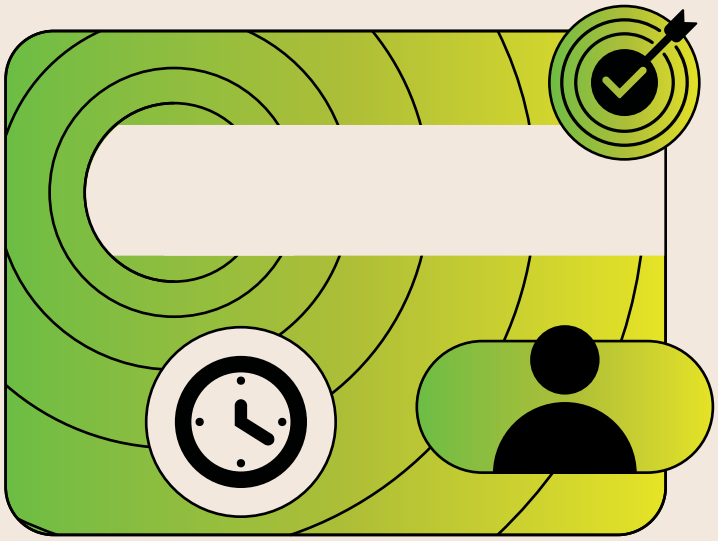
"Let's acknowledge that it's still early days with [GenAI] technology," Fuller said. "ChatGPT is a groundbreaking innovation, but it was released less than two years ago. We're just at the beginning of a generational change in technology."

States vary widely in their approach to AI, from "leaning in and embracing" to "holding back and waiting to see how it develops," Fuller observed. "We're actively figuring out how to optimize the AI stack to deliver helpful, predictable, useful, accurate, relevant, trusted outcomes."

DTS has added a dedicated AI team to support its IT organization. It's providing employee training statewide in both AI skills and responsible use. And it's experimenting with how best to deploy the technology in its operations.

One pilot project involves 250 employees across all Utah agencies who are trained on Google Workspace's AI capability.

"We're trying to give our employees superpowers," Fuller said. "We see them using GenAI tools to find more relevant answers to questions, to surface ideas, to help expand their thinking." Preliminary results from the pilot show



that 39% of users report saving one to two hours per week, and an additional 17% are saving more than two hours per week, according to Fuller.

A pilot program at the Utah Sales Tax Commission deploys RAG to help 200 call center agents better respond to constituents. The system connects a GenAI LLM to the state's tax law, agent training materials and previous call scripts to provide more accurate responses.

With nearly every vendor adding some kind of AI to their offerings, states will have plenty of tools to choose from, and RAG will help customize them. "We, as a state, will not be doing any

fundamental development of AI, like creating our own model. I don't think that's the role of a state," Fuller said.

Old-School AI

As exciting as GenAI is, there are times when more basic AI is the right tool. "I don't think we should expect GenAI to do things it's not good at," said Fuller. "Traditional AI is very good at analysis, pattern recognition and decision support." GenAI, not so much.

For example, Utah's Department of Transportation used AI to digitize decades worth of paper right-of-way maps and create a database of those rights-of-way, saving about 5,000 person-hours. Another AI-fueled database allows for quick identification of cattle brands — a task that used to involve flipping through giant books of pictures. A third reviews all the state's daily network traffic and flags anomalies that may be malware.

In 2025, AI Must Earn Trust

Although AI offers exciting possibilities, increasing use has brought some issues to the surface that agencies will contend with in the coming months.

"One of the biggest issues I see is the cost," Fuller said. "It's not cheap. And the large mega-vendors spending billions of dollars in R&D to build these tools are certainly hoping to get a positive return on investment. The licenses for the AI add-on to our productivity suite are more expensive than the productivity suite itself."

Getting your data ready for an AI model to use effectively can also be time consuming and costly. Training and familiarizing staff with its uses will be another significant expense.

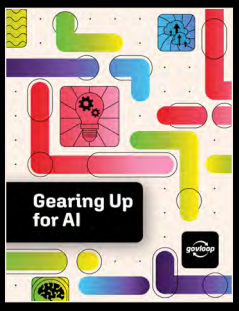
"The other big issue I see for AI tools in 2025 is trust," Fuller said. "The hallucination factor, the fact that these tools often give misinformation…that undercuts trust. We need to improve the models to the point of where they can be trusted," he said.

But Fuller believes these concerns can be worked out as AI matures and agencies grow more sophisticated in its use. "We do see it as a very significant and beneficial technology," he explained. "We're leaning in, and it's pretty exciting so far."

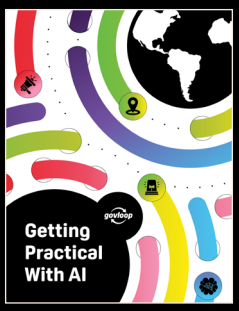


Conclusion

Our three 2024 AI guides have covered this fast-moving technology and how governments can make the most of it. If you've read this one, you'll want to read the earlier installments.



"Gearing Up for AI" focused on how agencies are getting ready for AI, including setting policy, understanding GenAI and making sure data is optimized for AI use. We interviewed Beth Simone Noveck, New Jersey's State Chief Artificial Intelligence Strategist and Director of The Governance Lab at Northeastern University, regarding how she sees AI changing government.



"Getting Practical With AI" showed how agencies are using AI now, with case studies and an in-depth interview with Alan Shark, author of "Artificial Intelligence: A Primer for State & Local Government," regarding how state and local governments can get up to speed — and even comfortable — with AI.

About GovLoop

GovLoop's mission is to inspire public-sector professionals by serving as the knowledge network for government. GovLoop connects more than 300,000 members, fostering cross-government collaboration, solving common problems and advancing government careers. GovLoop is headquartered in Washington, D.C., with a team of dedicated professionals who share a commitment to the public sector.

For more information about this report, please reach out to info@govloop.com.

Thank You

Thank you to AlphaSix, HPE and Pluralsight for their support of this valuable resource for public-sector professionals.

Authors

Lauren Walker, Senior Staff Writer
John Monroe, Director of Content
Candace Thorson, Managing Editor

Designers

Calista Lam, Graphic Designer
Andrew Blake, Motion Graphic Designer