AI FOR THE PUBLIC GOOD

A GOVLOOP GUIDE



Introduction

Much of the conversation around artificial intelligence (AI) right now focuses on the means: developing use cases, crafting policies, training the workforce. But what is the point of it all? That is, how will agencies use AI to improve their services to constituents, create new services and strengthen their stewardship of taxpayer dollars?

The **second installment** of our 2025 AI Guide series explores the emerging vision for how AI can serve the public good. We look abroad, highlighting AI initiatives in Estonia, the United Kingdom and other countries that could serve as models for agencies in the United States.

But using AI for the public good also requires good governance: As the number of AI initiatives increases, how can agencies prioritize and consistently deliver high-quality work? With that in mind, we spoke with Zach Whitman, Chief AI Officer and Chief Data Officer at the U.S. General Services Administration, about how GSA has created an AI governance infrastructure to guide its initiatives.

Finally, this guide includes a bonus feature: an AI glossary. AI has become such a hot topic of conversation, but people often do not take the time to define the terms they use. Our glossary can help fill in some common gaps. Definitely a clip-and-save.

This all sets the stage for our third and final AI Guide of the year, which will come out in November. In it, we distill the lessons learned about AI's role in government in 2025 and highlight key issues to address in the year ahead.

Contents

3 **How to Speak AI: A Glossary for the Rest of Us**

5 **How to Minimize the Risk on the Road to AI**

6 **AI for the Public Good: A Global View**

7 **Case Study #1: AI Helps the U.K. Government Make Sense of Public Input**

8 **Case Study #2: Estonian Parliament Enlists AI to Support Lawmakers**

9 **Building a Network for AI Cybersecurity**

10 **Case Study #3: Singapore Puts AI to Work on Municipal Problems**

11 **Case Study #4: New Zealand Offers Employees an AI-Intensive Course**

12 **How to Move Ahead on Your AI Journey**

13 **Government Spotlight: GSA Maps a Journey to an AI-Optimized Future**

15 **Conclusion + Acknowledgments**

How to Speak AI: A Glossary for the Rest of Us

We've all been there: Reading an article or watching a webinar about artificial intelligence (AI), we come across a term that we know we've heard and think we understand but...we're not quite sure. With that in mind, here are some common AI terms you might want to know in the months and years ahead.

AGENTIC AI: Software that, given a goal, can make decisions and act with limited supervision, making it possible to automate more comprehensive processes and not just individual tasks.

BIAS: Systematic errors in AI output that discriminate against some groups (e.g., a facial recognition system that was trained on people with light-colored skin so it cannot accurately identify individuals with dark skin).

EMERGENT BEHAVIOR: Behavior or capabilities that a given system was not specifically designed to produce (for better or worse).

EXPERT SYSTEM: A system designed to replicate human decision-making in a given field (e.g., chemistry), building on both an extensive domain-specific knowledge base and clear rules of logic.

EXPLAINABILITY: The ability to understand how a model "thinks," meaning how it reaches a decision or outcome.

GENERATIVE AI: Software that can be used to create similar but new content. GenAI typically is based on a **large language model** (LLM), software that has been trained to process human language by analyzing a vast amount of text or other content. Alternatively, GenAI might be based on a **small language model**, which involves less training data, often because it is focused on a specific domain of information (e.g., medical content).

MACHINE LEARNING: Involves using data to train algorithms, which then can analyze real-world data and find patterns, make predictions and so on. The algorithms can be trained with human oversight and guidance (**supervised learning**) or without it (**unsupervised learning**).

NEURAL NETWORK: A type of ML that uses multiple, interconnected nodes that process data in collaboration, much like the neurons in a human brain. **Deep learning** uses a multi-layered neural network to tackle very complex tasks, such as image or speech recognition or autonomous operations. (LLMs are developed using deep learning techniques.)

PROMPT ENGINEERING: The process of writing a question or series of questions (or similar inputs) to provide GenAI with enough guidance to produce the outcome you want.

RESPONSIBLE AI: The practice of developing AI models and systems that minimize risks to individuals and society at large by addressing concerns around fairness, privacy and security, and transparency or explainability.

RETRIEVAL-AUGMENTED GENERATION (RAG): A process for supplementing the output of a general LLM with more context-specific information from a reliable source of information specific to a given environment (such as an agency's knowledge base) or domain.

TURING TEST: A test of whether a system, such as a chatbot, has achieved human-like intelligence. A system passes the test when a person interacting with it online thinks it's a human.

Four Types of AI Functionality

Not every AI system has the same functionality, or core capability. Today's applications generally fall into one of two categories:

Reactive Machine AI: Although they can analyze large volumes of data, these systems have no memory. That means they can't track how well a decision or action worked out and use that information to improve outcomes. For example, a chess application based on reactive machine AI will not learn from game to game.

Limited Memory AI: Such a system has some memory so it can change how it makes decisions based on experience. Although it can't retain data in the long term, its algorithms do improve over time. GenAI and autonomous cars both use limited memory AI.

Two more types of AI are possible but haven't been developed yet:

Theory of Mind AI: Such a system would be able to understand how people (and other creatures) think and feel and then interact with them accordingly.

Self-aware AI: Such a system would have human-like consciousness and be capable of its own thoughts and emotions. Not everyone believes this is even feasible (or desirable).

The Evolution of AI

Although AI has been around for decades, the technology is very much in its early stages. Experts expect that AI applications eventually will fall into one of three general tiers, based on the range of their capabilities:

TYPE	DEFINITION	SAMPLE USE CASES
Narrow AI	Also called weak AI, it's a type of AI designed to carry out a specific task or set of tasks or to solve specific types of problems.	Voice assistants, facial recognition systems, disease mapping
Strong AI	Also called Artificial General Intelligence (AGI) — and at this point, still just theoretical — it will be able to apply human-like reasoning to a wide range of tasks or domains, with the ability to learn as it goes.	In a December 2023 report, the Government Accountability Office highlighted Lieutenant Commander Data, the android from the sci-fi series "Star Trek: The Next Generation" TV show, as an example of AGI.
Super AI	Also called Artificial Superintelligence — and at this point also theoretical — it will outperform humans in every way, including cognitive ability, creativity and even emotional intelligence.	TBD!

Industry Spotlight

How to Minimize the Risk on the Road to AI



Agencies are quickly realizing they cannot achieve the kind of efficiency that's expected of them without using AI. Yet they cannot move ahead with AI unless they learn to manage the risks associated with it. And failing to move ahead is not an option.

"The key is to reframe the challenge," says Chad Cisco, Chief Customer Officer at DataRobot. "Risk shouldn't be viewed as a roadblock to efficiency but as a manageable part of the process. With the right strategies and tools, those risks can be mitigated." This shift in perspective is now shaping how the federal government approaches AI implementation.

In this video interview, Cisco discusses how agencies can accelerate their adoption of AI without putting their systems and data at risk.

Topics include:

- The key security risks with AI and how to address them
- The essential components of an AI governance and ethics strategy
- Data privacy concerns specific to AI

"OMB recently put out a memo that asked chief data and AI officers to think about themselves as enablers for using AI. If you've thought about these risks ahead of time and established your frameworks, you're going to be able to move faster."

– Chad Cisco, DataRobot

About DataRobot

DataRobot delivers AI that maximizes impact and minimizes business risk. Its AI applications and platform integrate into core business processes so teams can develop, deliver, and govern AI at scale. DataRobot empowers practitioners to deliver predictive and generative AI and enables leaders to secure their AI assets. Organizations worldwide rely on DataRobot for AI that makes sense for their business — today and in the future.

For more information, visit the DataRobot website and connect with them on LinkedIn.

DataRobot

AI for the Public Good: A Global View

The United States is not alone in its efforts to incorporate AI into government operations and services — what the Organisation for Economic Co-operation and Development calls the machinery of government. In a 2024 report, OECD identified three main ways AI can improve public administration:

- Increase productivity with more efficient internal operations and more effective public policies
- Make the design and delivery of public policies and services more inclusive and responsive to the evolving needs of citizens and specific communities
- Strengthen governments' accountability by enhancing their capacity for oversight and supporting independent oversight institutions

In the following pages, we look at examples of AI at work, with an emphasis on efficiency and effectiveness, which are guiding efforts here in the United States.

AI Case Study #1

AI Helps the U.K. Government Make Sense of Public Input

What's Happening

For policymakers in the United Kingdom, public comments are essential to the rulemaking process, but often they can feel like too much of a good thing. Now, an AI tool dubbed Consult is making it easier and less expensive to review and summarize that input.

How It Works

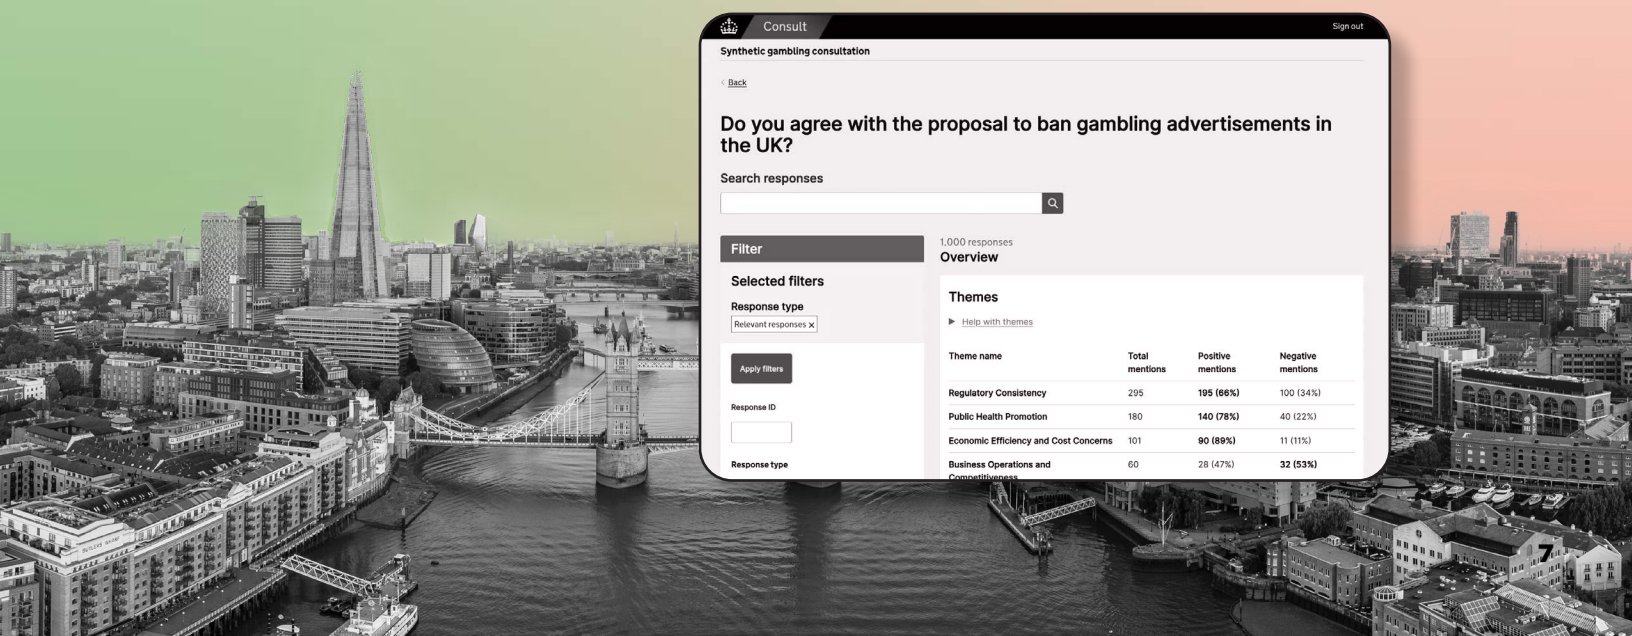
U.K. officials seek public comments on proposed policies (what they call "consultations") about 600 times a year. Although some receive a manageable volume of feedback, the response to others can be overwhelming, numbering in the tens of thousands. It takes about three months for a team of 25 analysts to analyze 30,000 comments, and some consultations receive twice that.

Consult, developed by the U.K. government's Incubator for Artificial Intelligence (i.AI), takes a two-step approach to reviewing comments. First, it identifies key themes, providing a basic description and context. Then it shows how those themes map back to specific responses. The result is a dashboard that policymakers can use to explore the comments.

In its first real-life use, the Scottish government used Consult to review feedback on ideas for regulating non-surgical cosmetic procedures. The results of Consult's analysis closely matched those of a review by human analysts, according to i.AI.

Why It Matters

This is an archetypal AI use case: By using AI to address a significant pain point, the government expects to save millions of pounds per year in staff and contracting costs, while freeing humans to focus on higher-level analysis. "No one should be wasting time on something AI can do quicker and better, let alone wasting millions of taxpayer pounds on outsourcing such work to contractors," said Technology Secretary Peter Kyle in a statement.



AI Case Study #2

Estonian Parliament Enlists AI to Support Lawmakers

What's Happening

Estonia's parliament, called the Riigikogu, has developed an AI tool to help its lawmakers and their staff conduct research in the legislative database, known as EMS, and in the Riigi Teataja, the country's public record of laws.

How It Works

When drafting bills, members of parliament (MPs) or their staff frequently need to refer to existing laws and associated legal documents. Although all that information is available in the two central databases, finding and analyzing the documents can be time-consuming.

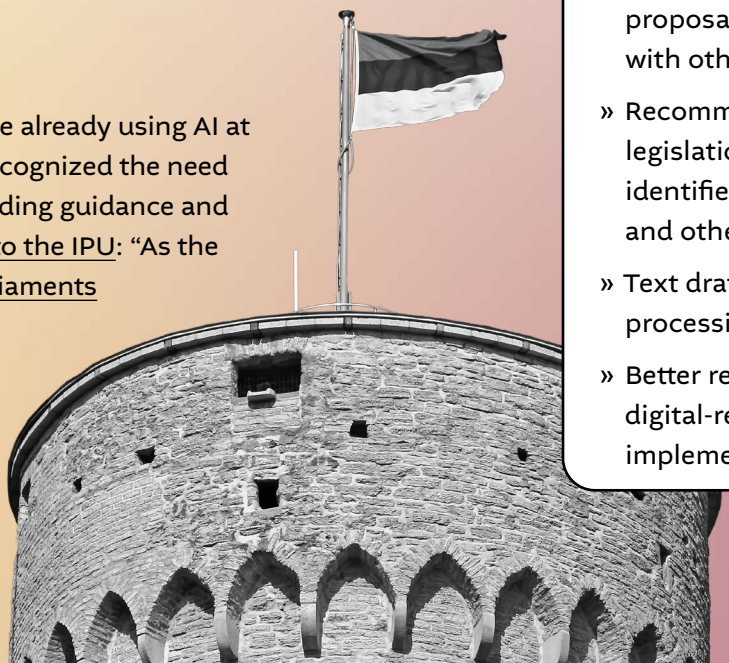
The new tool, currently in pilot phase, provides prompting tools that help users find specific texts, analyze their content and, in the case of active proposals, clarify their status and the status of any amendments. The parliament first rolled out the tool to its legal and research department, which supports MPs.

The Riigikogu recently began training MPs and staff on the new tool as part of a broader training program on the effective use of AI. That training is based on guidelines developed by the Inter-Parliamentary Union, an association serving parliamentary bodies worldwide.

This is not the Riigikogu's only foray into AI. It also uses AI to transcribe parliamentary sessions. Transcripts are now available in about 20 minutes, rather than in hours, with only a 5% error rate.

Why It Matters

Given that so many people are already using AI at home and at work, Estonia recognized the need to put more energy into providing guidance and practical training, according to the IPU: "As the IPU's Guidelines for AI in Parliaments make clear, having informed users is a key part of the equation for promoting the safe and effective use of trusted AI tools."



Parliamentary AI-based Applications

Here are some ways that legislative bodies can apply AI to their work, according to a study by the Westminster Foundation for Democracy:

- Real-time subtitling of MPs' speeches
- Reliable voting systems in plenary and committees
- Generation of content for speeches and written questions
- Support in information retrieval
- Examination of legislative proposals for interactions with other regulations
- Recommendations on legislation based on identified gaps, problems and other relevant laws
- Text drafts for further processing
- Better regulation and digital-ready policy implementation

Industry Spotlight

Building a Network for AI Cybersecurity



Artificial intelligence (AI) is already making work easier in countless ways, and it has the potential to transform how agencies function. In the cybersecurity realm, it can detect threats and take action faster and more accurately than ever before. But making the most of AI requires a high-performance network infrastructure that legacy systems can't provide.

Upgrading network infrastructure is more than just fixing what's broken. Done right, it can integrate AI so agencies get its full benefits, help them address new and growing security challenges, and automate more functions. But modernization isn't easy. Technologies such as Cisco's networking as a service (NaaS) can help.

In this video interview, Kevin Manwiller, Operations Director for National Security at Cisco Systems, discusses how AI can transform agency security and the network structure that's needed to support it. Topics include:

- Why agencies need AI to counter the current, and future, threat landscape
- What network capabilities best support AI cybersecurity
- How AI agents will change security response

"If you have the greatest AI tool in the world, but the network is lagging behind it and not providing the type of performance it needs, the AI tool is going to be viewed as a failure. Without the right modern network, AI actually can become a liability, instead of an asset."

— Kevin Manwiller, Cisco Systems

About Cisco

Cisco is the worldwide technology leader that securely connects everything to make anything possible. Its purpose is to power an inclusive future for all by helping customers reimagine their applications, power hybrid work, secure their enterprise, transform their infrastructure, and meet their sustainability goals.

Learn more about Cisco

AI Case Study #3

Singapore Puts AI to Work on Municipal Problems

What's Happening

Twenty-five years ago, many local governments took their initial steps into e-government by offering online applications that let constituents report potholes, broken streetlights, litter and other nuisances. Singapore is now rebooting that vision with an AI-powered tool called OneService Kaki Chatbot.

How It Works

Before developing the chatbot, Singapore created a OneService app based on old-school internet technology. To report a problem, users chose from a list of possible issues (pothole, streetlights, etc.) and then filled in various fields.

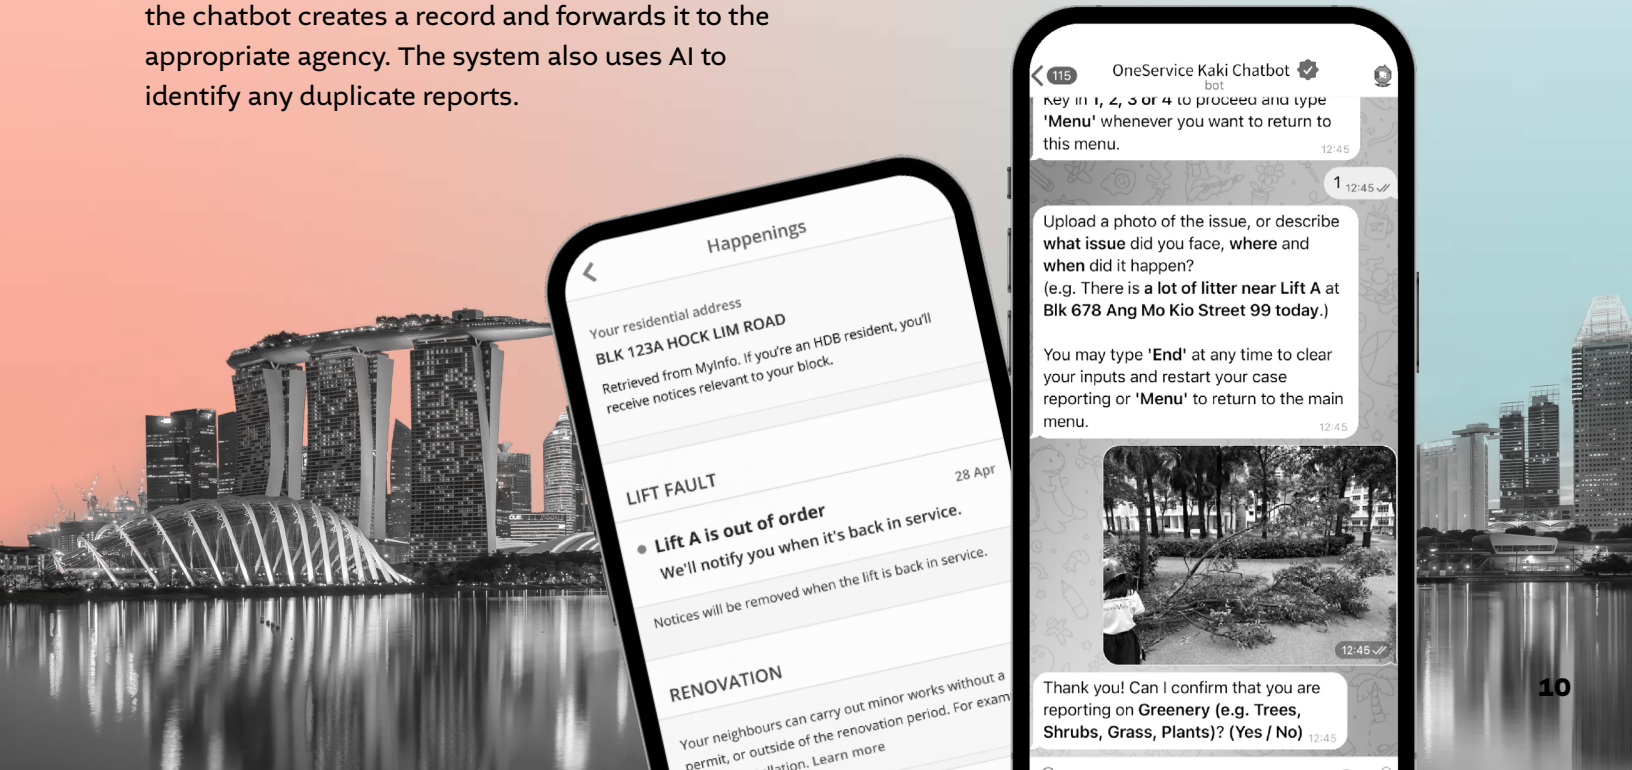
The new chatbot provides a more intuitive interface. Rather than working with pull-down menus or form fields, a user simply describes the problem in their own words. Then the chatbot asks follow-up questions to get more information and requests photos. The net result often is a more comprehensive report than was possible through the traditional app.

Using that information, plus geo-location data, the chatbot creates a record and forwards it to the appropriate agency. The system also uses AI to identify any duplicate reports.

Why It Matters

The result is a better experience for both constituents and the government employees serving them, as AI takes a lot of drudge work out of the process. "With this, citizens can easily provide municipal feedback without having to figure out which agency they should contact," Singapore's 2023 National AI Strategy states. "Agencies can also attend to cases more quickly."

That chatbot, which handles about 30,000 cases per month, has reduced response times by up to two working days, according to Singapore's Government Technology Agency. Constituents can access the tool through two popular social media apps, WhatsApp and Telegram.



AI Case Study #4

New Zealand Offers Employees an AI-Intensive Course

What's Happening

Looking to spur AI adoption across the country's public sector, New Zealand's central government has developed a training program designed to get public-sector employees up to speed on AI basics.

How It Works

The AI Foundational Development Programme, open to all employees, is a five-week upskilling course in which participants not only learn key concepts but also get hands-on practice developing use cases. Topics include creating value with AI, implementing and evaluating AI programs, and, of course, using AI safely and responsibly (see sidebar).

The course concludes with capstone projects in which participants partner with agency colleagues to develop a use case and pitch it to a panel of senior government leaders. The first cohort, which ran in May and June 2025, included 50 participants from 10 agencies.

The three-part AI Masterclass series is designed to help senior executives understand core AI concepts and key decision-making considerations and to think about how to apply them in government.

A Five-Step Approach to Learning AI

Module 1: Creating Value with AI
The basics of AI, its implications for government, and how it can drive productivity and growth

Module 2: Safe and Responsible AI
Exploring AI through a public service lens, focusing on ethics, privacy and security

Module 3: Implementing AI
The real-world impacts of AI on people and processes, including the trade-offs between building vs. buying AI solutions

Module 4: Evaluating AI
Applying the basics to AI use cases

Module 5: Capstone Projects
Developing and presenting a practical use case project to a panel of senior government leaders

Why It Matters

The two courses, developed by Government Chief Digital Officer Paul James in partnership with the Public Service Commission's Leadership Development Centre, look to help agencies accelerate their adoption of AI without driving up risks.

"Across the public sector, AI offers a real opportunity to improve decision-making and deliver better outcomes for New Zealanders," said Judith Collins, Minister for Digitizing Government. "I want public service leaders to embrace AI and support their teams to use it safely and effectively."



Industry Spotlight

How to Move Ahead on Your AI Journey



As agencies explore a growing range of use cases for artificial intelligence (AI), they face a looming challenge: Their existing IT operations, while perhaps sufficient for traditional applications, are not ready to meet the demands of full-fledged AI initiatives. Replacing a horse-and-buggy with a sports car won't pay off if you're still driving on a dirt trail.

As always, it comes back to people, process and technology. The task begins with understanding where you're trying to go, that is, how you see AI augmenting your workforce. Once you've defined that use case, you need to look at building an AI model and feeding it with the necessary data. And then it all comes down to the infrastructure. Agencies need to lay a solid foundation for AI, building a technology stack that can keep pace with their evolving requirements, said Stephen Tuomey, a Data Center Account Executive at Dell Technologies, which partners with SHI on AI projects.

In this video interview, Tuomey provides a roadmap that can help agencies accelerate their AI journey. Topics include:

- Selecting the right AI model for a given use case
- Building the technology stack for testing and deploying that model
- Scaling the technology stack as the model evolves

"It's about getting AI into the data center and updating that data center from where it is today, with maybe some legacy hardware and old-school virtualization, and moving it forward into a cloud-native architecture."

— Stephen Tuomey, Dell Technologies

About SHI and Dell Technologies

As a Dell Federal Gold partner, SHI provides innovative IT solutions that enhance federal government operations. Together, SHI and Dell deliver cutting-edge AI and hyper-converged infrastructure, reshaping the IT landscape to drive operational efficiency and decision-making.

Learn more about SHI and Dell Technologies

SHI | PUBLIC SECTOR

DELL Technologies FEDERAL GOLD PARTNER

Government Spotlight

GSA Maps a Journey to an AI-Optimized Future

Part of what makes GenAI so powerful is that, unlike many cutting-edge technologies, just about anyone can use it without having great technical expertise. But to realize AI's full benefits, agencies need to take a cohesive approach to deploying, managing and monitoring the technology.

Because here's the catch: **Although anyone can use GenAI, they don't necessarily know how to use it effectively and safely.** That's the idea behind a new initiative at the U.S. General Services Administration, said Zach Whitman, GSA's Chief AI Officer and Chief Data Officer.



"When we started this GenAI adoption effort, we were thinking, 'What could be more important than making sure that our folks have the tools, are fully trained on how these tools work and then start to incorporate them into their day-to-day workflow?'" Whitman said.

The result of those efforts, while geared toward its own workforce, could soon extend to other agencies, he said.

Something for Everyone

The centerpiece of GSA's initiative is GSAi, a GenAI platform modeled after commercial offerings but custom-built to address public-sector concerns around security and reliability. The platform includes three components: a chatbot, an application programming interface that integrates GSAi with other systems and a monitoring console.

Agency leaders are especially interested in seeing how GenAI can help employees reduce the drudgery in their daily work. For example, contracting specialists have used the tool to check IT for compliance with Section 508 accessibility requirements. That's an important task that GenAI can handle easily.

"We wanted to make sure that folks reviewing these offerings had time to focus on the real meat of their work, rather than having to say, 'Oh, you forgot to add this compliance thing,'" Whitman said. That "gives our folks time to focus on things that need more focus."

The monitoring console is especially important because it enables the agency to assess the large language model's behavior and performance and ensure it meets users' needs — or to fine-tune it as appropriate.

Since launching GSAi in March 2025, GSA employees have adopted it relatively rapidly. And Whitman also has heard from several agencies looking at the same approach.

Although each agency has its own requirements, he said, many employees do similar back-office activities, such as drafting memos, preparing presentations and conducting research. GSA believes its platform could become the basis of a shared services architecture for agencies across the federal government.

"A shared platform is one way to consolidate the modernization effort in a single place, rather than us all having to reproduce the wheel independently of one another," Whitman said.

Doubling Down on Governance

Agency leaders are encouraging employees to work with the tool and test use cases that align both with GSA goals and policies.

The governance model has two levels:

- **A board of executives** sets the overall strategic direction in terms of identifying the kind of use cases the agency should pursue. For example, should they focus primarily on back-office or higher-visibility work? What is the agency's risk appetite?
- **A safety team**, which Whitman describes as the "muscle" behind the governance board, evaluates individual use cases, looking at technical, legal, ethical and privacy issues, among others. Their first task is to determine if a given project is appropriate for AI.

Because that discussion goes beyond the purview of technical specialists, the board includes subject-matter experts with a deep understanding of GSA business, Whitman said.

"With this very diverse cadre of expertise, we're able to come to a decision as a group, to say, 'Yes, this is an appropriate use of technology,'" he said. At that point, the proposal can go through the technology approval process.

Looking Beyond Chatbots

Right now, the goal is to turn this kind of platform into a commodity, Whitman said, one where its use becomes the norm. But that's just a short-term objective.

Whitman said he expects to see the chatbot become less important over time. First, agencies will integrate GenAI into larger systems and workflows, improving the efficiency not just of employees, but of the workflows themselves. Next, agencies will look to optimize those workflows with AI in mind, driving even greater efficiencies.

"I think that these types of workflows and pattern changes are the things that we need to keep an eye on for future product design, future workflow design," Whitman said.



Conclusion

This guide is the latest look at the rapid revolution of AI in the public sector. Be sure to check out articles from previous guides:

March 2025 — Government Gets Down to Business With AI

» Four case studies in driving better outcomes

» What Alameda County, California, has learned about getting employee buy-in

November 2024 — AI: Where We Are, Where We're Going

» On the road to responsible AI

» How Utah is mapping a course to AI success

July 2024 — Getting Practical With AI

» A look at six common use cases taking hold in the public sector

» An interview with Alan Shark, author of "Artificial Intelligence: A Primer for State & Local Government"

» Pre-register for the next guide in our 2025 AI Series

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.

govloop.com | @govloop

Thank You

Thank you to Cisco, DataRobot, Dell Technologies and SHI for their support of this valuable resource for public-sector professionals.

Authors

John Monroe, Director of Content

Candace Thorson, Managing Editor

Lauren Walker, Senior Staff Writer

Designer

Kaitlyn Baker, Senior Creative Manager