

# Gearing Up for AI



## Introduction


The modern history of artificial intelligence (AI) goes back to 1950, when Alan Turing published his paper, Computing Machinery and Intelligence, raising the question, “Can machines think?” In 1956, John McCarthy, Marvin Minsky, Nathaniel Rochester and Claude Shannon introduced the term “artificial intelligence.” Arthur Samuel popularized the term “machine learning” in his 1959 paper about a program that could play checkers. Five years later, one of the first chatbots to use natural language — ELIZA, a kind of virtual psychotherapist, arrived. (Try it out.)

It wasn’t until 1997 that IBM’s Deep Blue chess-playing AI was able to beat human grandmaster Garry Kasparov, who was at the time the top chess player in the world. (Kasparov rallied, and bested the machine in the six-game match.)

Since then, AI’s advance has been fast and furious. Governments have responded with regulation and legislation — and also by adopting these tools to enhance customer experience, data-based decision-making and resource management, among other uses. In this guide, we’ll show where AI has been and where it may be going in 2024.

## Table of Contents


- 3** **Timeline of Federal and State AI Guidelines + Legislation**
- 6** **Managing Your Data for AI**  
with Pure Storage
- 7** **Leveraging AI to Speed Incident Detection and Response**  
with Elastic
- 8** **AI in 2024: Laying the Groundwork**
- 11** **How MLOps Helps Agencies Get the Most Out of AI**  
with Red Hat
- 12** **Getting Comfortable With AI**
- 13** **Meeting AI Where You Are**  
with Four Inc. and IBM



# Timeline of Federal and State AI Guidelines + Legislation

## ▲ FEDERAL



## ✦ STATE



2016

Obama administration releases Preparing for the Future of AI, the first federal AI guidance. ▲


The National Science and Technology Council releases The National Artificial Intelligence Research and Development Strategic Plan. ▲



2018

Congress establishes the National Security Commission on Artificial Intelligence. ▲


**CA:** Enacts the California Consumer Privacy Act.  
**VT:** Creates an AI task force. ✦



2019

The Office of Management and Budget (OMB) releases draft guidance for regulation of AI applications. ▲

**CA:** Requires algorithms to be evaluated for bias in pre-trial decision-making.  
**IL:** Requires employers using AI to notify applicants.  
**NY:** Creates a commission to study regulation of AI, robotics and automation. ✦




2020

H.R. 6216, the National Artificial Intelligence Initiative Act of 2020, passes as part of National Defense Appropriations Act.

OMB releases final version of guidance for regulation of AI applications.

**MD:** Prohibits use of facial recognition without applicant's consent.

**NYC:** Limits use of facial recognition in hiring.



2021


The National Security Commission on Artificial Intelligence releases its final report.

**CO:** Enacts Colorado Privacy Act and an act to restrict insurers' use of external consumer data.

**IL:** Requires video employment interviews to be evaluated for bias.

**VA:** Enacts Consumer Data Protection Act.

**WA:** Creates a work group on fairness of automated decision-making systems.



2022

Biden administration releases its Blueprint for an AI Bill of Rights.

**CT:** Gives consumers rights over their data.

**VT:** Regulates the state's use and oversight of AI.

## 2023

The National Institute for Standards and Technology (NIST) releases its Artificial Intelligence Risk Management Framework. ▲

NIST launches the Trustworthy and Responsible AI Resource Center. ▲

President Biden releases the Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence. ▲

OMB introduces draft guidance on implementing the EO. ▲

**CA:** Requires an inventory of all automated decision systems used by the state.

**CT:** Establishes an Office of Artificial Intelligence to study and regulate use of AI.

**DE:** Enacts Delaware Personal Data Privacy Act.

**IA:** Enacts Consumer Data Protection Act.

**IL:** Establishes a generative AI (GenAI) and natural language processing task force.

**IN:** Adds consumer data protection to state trade regulations.

**LA:** Calls for study and report on impact of AI.

**MT:** Enacts Montana Consumer Data Privacy Act.

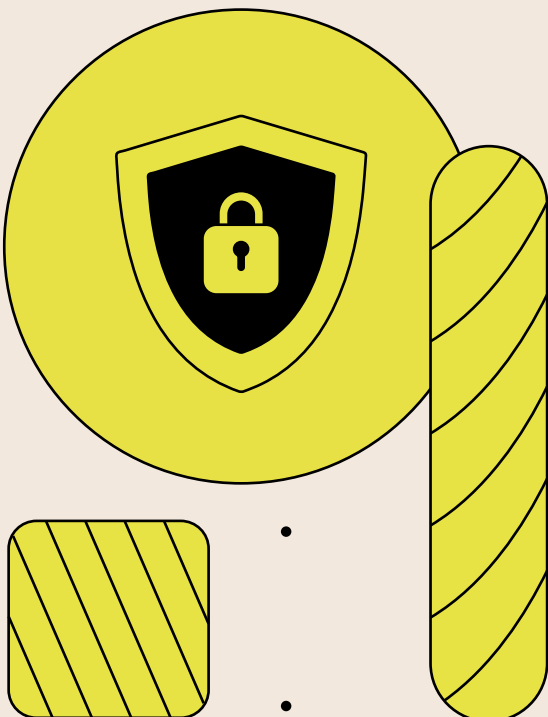
**NY:** Creates a commission to study how to regulate AI.

**OR:** Enacts Oregon Consumer Data Protection Act.

**TN:** Enacts Tennessee Information Protection Act.

**TX:** Creates an AI advisory council.

**TX:** Enacts Texas Data Privacy and Security Act. ✦



# Managing Your Data for AI

An interview with Richard Barlow, Principal Technologist for State Local Government and Education (East), Pure Storage

Agencies struggle with data for many reasons. It grows amazingly fast and draws from novel sources, including Internet of Things devices. Finding qualified employees to handle it, especially in the public sector, can be difficult. And there are compliance and cybersecurity concerns, and technology integration challenges.

Mismanaged data can lead to poor decision-making, increased risk and even legal fallout, but the No. 1 consequence is loss of trust, said Richard Barlow with Pure Storage, which provides agencies with fast, secure, energy-efficient data storage. The public remembers your mistakes.

“You can have a beautiful white shirt, flawless. But if there’s one tiny stain on the collar, everybody will see that one stain,” he said. “They won’t see the rest of the shirt at all.... All they remember is the one thing you or the agency did wrong. It’s never the 50 years of flawless service.”

## Recipe for AI Analytics

AI helps agencies make better, safer use of data, but there’s a recipe for using AI-driven analytics effectively, Barlow said. Agencies need enormous amounts of data, potentially millions of data points. They need the right data, which could be years-old information that once seemed irrelevant.

“It used to be [that] you go through your data, you find the important pieces and you throw everything else away,” Barlow remembered. “With AI, you can never do that.”

And agencies must use appropriate AI models. “As you grow more and more data, you have to decide when [it’s] time to start following the model, when you trust it,” he said. AI models sometimes produce counterintuitive results, but that’s OK, he said, as long as enough of the right data was factored in.

## Exciting Possibilities

AI tools can seem almost magical, but Barlow said natural language processing is one of the most exciting in the public sphere. “You can tell a system what you want in normal English ... and instead of having a data scientist sit down and create a complex model, the system, using your verbiage, creates it for you,” he explained.

ML models detect fraud by comparing data — for instance, your tax return against tax filings that millions of similar individuals submitted.

And geospatial analytics, Barlow said, will improve disaster response planning, weather forecasting and municipal development, among other possibilities.

## Getting AI-Ready

However, agencies need three things before launching AI initiatives, according to Barlow: a flexible mindset, an overarching AI strategy and a data management plan. Pure Storage provides software and hardware to help agencies maximize data’s potential, and it reduces the complexity and cost of managing IT infrastructure, he said.

The FBI, CIA, Department of the Navy and other agencies with immense data needs rely on Pure Storage. Barlow noted the company’s impressive customer experience ranking: an industry-leading 82.4 Net Promoter Score. “If you have a problem, we’re there with you and we’ll help you fix it,” he said.



# Leveraging AI to Speed Incident Detection and Response

An interview with Nathan Stacey, Senior Manager, Solution Architects, at Elastic

With greater observability, federal agencies can gain actionable insight into the performance of their IT systems and applications. And by adding AI to their IT operations (AIOps), they can supercharge their ability to apply intelligent incident detection and response in order to speed remediation and keep the mission on track, said Nathan Stacey, Senior Manager with Elastic.

## AIOps: Improving Observability, Streamlining Operations

“Observability is the collection and coordination of something’s health via data,” Stacey said. “AIOps is an always-improving OODA [observe, orient, decide, act] loop built from that observability. Just like turning the lights on in a room, the more data (light) and speed to search that data, the better the decision.”

App owners have astronomically more levers than ever to pull for better performance, cost, improvements and reconfigurations, Stacey noted. AIOps maximizes the value of these levers by having them pulled automatically. And once agencies tackle observability through the application lens, it’s time to turn attention to full-stack observability, he said. When agencies can generate insights throughout their IT ecosystems and cloud environments — public, hybrid, on premises and multi-cloud — they can minimize potential downtime.

Although not every agency will want to achieve full-blown observability immediately, most will at least be moving in that direction. A unified platform can support the effort, driving cost savings and tool consolidation.

And a unified platform offers agencies a single means to understand operational and mission data, context, and correlation across telemetries by making sure all systems speak the same language, Stacey said.

## Detecting Hard-to-Find Problems, Supporting Collaboration

“Observability automatically combines different log fields so when there is a problem, we know what logs matter,” Stacey explained. AIOps looks at all of those logs, sometimes millions of them, and identifies anomalies related to the exact problem. “Where observability helps us see the needle in the haystack, AIOps removes the hay and hands us only the needles.”

With the same AIOps workflows at desks, in the field or at central operations, teams can collaborate and troubleshoot in real time. Each person needs unique access and data control to work together to resolve issues quickly; Elastic provides this role-based access, so data owners can share data with anyone they want, no matter their location.

“Elastic is built for common users trying to improve their missions and is easy to learn and configure,” he said. “Our speed and scale allow for limitless usage of observability data, and our flexibility allows limitless mission needs to matter inside the IT world, enabling users to bring as much IT performance as possible to the mission.”

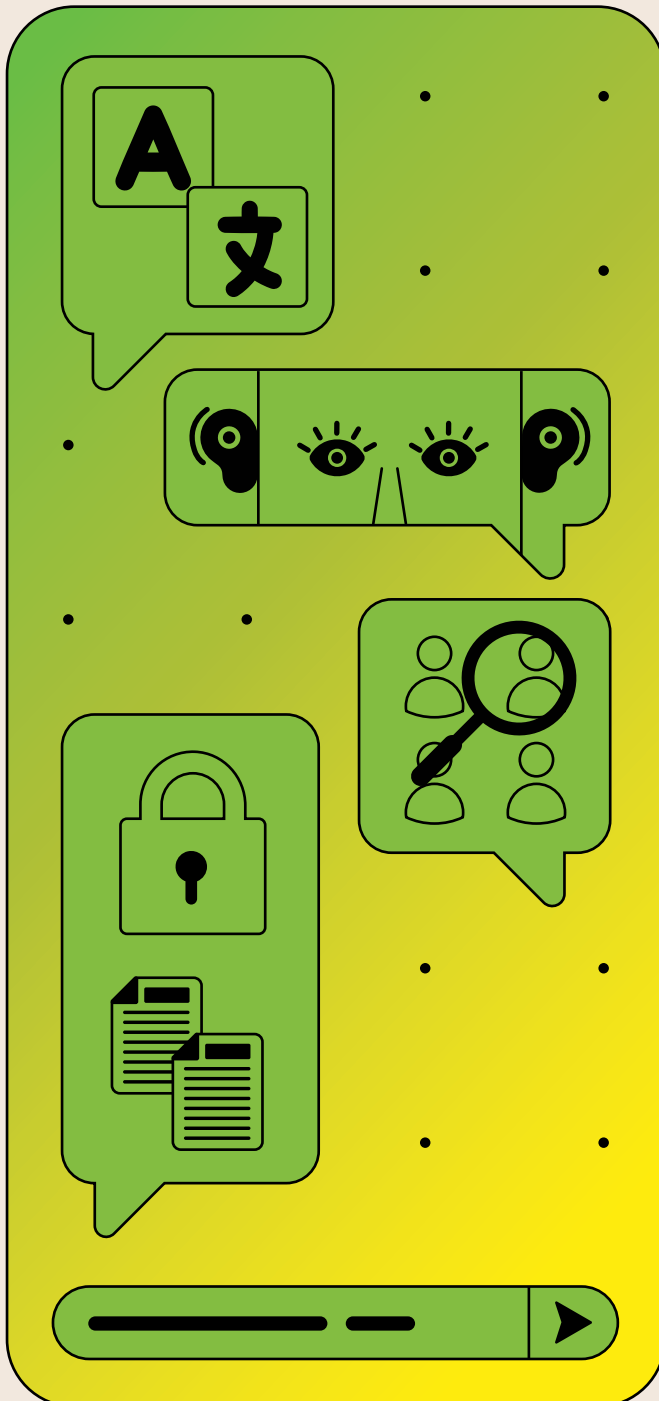
Partnering with AWS, Elastic delivers search-powered solutions that help people find what they need faster while keeping applications running smoothly, securely and more productively, he added. That promotes system resiliency, as agencies that embrace solutions incorporating AIOps ultimately benefit from intelligent automations and domain-specific ML rules. And organizations with cloud monitoring solutions, Stacey said, gain real-time insights into complex hybrid and multi-cloud environments.





## AI in 2024: Laying the Groundwork

Last year, AI stepped into the spotlight. This year, agencies will define its role. Here are four considerations as you plan your next steps.



### 1. Understand Generative AI's Potential and Pitfalls

The public release of GenAI tools — chatbots such as [ChatGPT](#) and image generators such as [DALL-E](#) — late in 2022 drew the world's attention to AI. By January 2023, ChatGPT had [100 million monthly active users](#), making it the fastest-growing consumer software application in history.

It quickly became clear that GenAI offers great possibilities, but also significant risks.

GenAI can improve government in multiple ways, including:

- Translating language and converting documents and videos for users with hearing and vision impairments
- Identifying groups that can benefit from more outreach
- Summarizing meetings, reports and other documents
- Analyzing public feedback
- Converting legacy software code to modern programming languages and streamlining new code
- Improving cybersecurity
- Optimizing resource allocation and energy efficiency

The primary risk of using online versions of GenAI tools is that they are public platforms. You can't assume privacy or accuracy, as the U.S. Senate noted in its [guidelines](#) last December.

Nearly as dangerous are GenAI's "hallucinations," — its well-publicized tendency to put words together in a way that sounds plausible but isn't accurate. Every AI output must be double-checked against other resources and reviewed by a human.

## 2. Cultivate AI Expertise on Staff

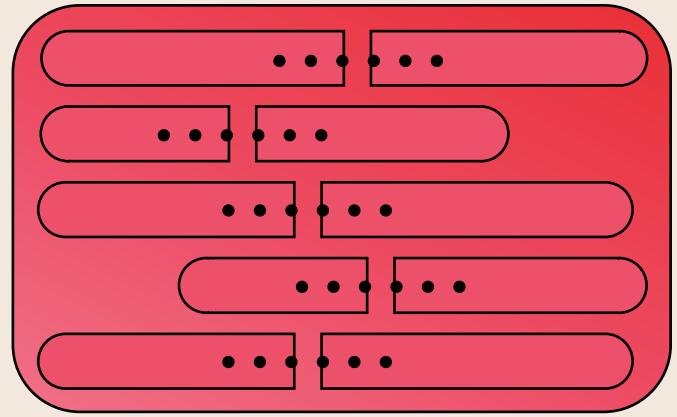

Federal AI guidance emphasizes the need for AI skills in the workforce — in more than just technical roles. OMB advises agencies to put AI-savvy people in both mission and program offices, including designers, behavioral scientists, contracting officials, managers and attorneys.

It turns out that many of the skills required to work effectively with AI are “soft” skills. The list of AI competencies from the Office of Personnel Management includes “creativity and innovation,” “integrity” and “political savvy.”

A Government Accountability Office study regarding the Department of Defense recommended answering these questions upfront:

- Who is included in the AI workforce?
- Who should be included in the AI workforce?
- Which positions require personnel with AI skills?
- What is the current state of your AI workforce?
- What are your future requirements?

Ongoing learning will be an essential element of life with AI. Agencies must help their existing employees adapt, both by offering training in AI-related skills and mapping career pathways for people whose jobs are substantially altered by AI adoption.



## 3. Put AI to Work for Cybersecurity

One of the most promising arenas for AI is cybersecurity. As the number and variety of cyber threats have grown, it’s become impossible to keep up manually. Even traditional software systems can’t keep pace, according to the IEEE Computer Society.

Agencies can use AI’s strengths in pattern recognition and data analysis to protect systems and data. AI can:

- Spot anomalies, such as suspicious login attempts, in real time
- Automate incident response, reducing the time from incursion to remedy
- Predict likely threats, by both assessing system weaknesses and staying up to date on evolving attack methods

In August 2023, the Biden administration launched its Artificial Intelligence Cyber Challenge. Competitors will use AI to identify and fix vulnerabilities in some of the country’s most essential software infrastructure, such as code that runs the internet, the electric power grid and transportation systems. The Defense Advanced Research Projects Agency is partnering with AI vendors including Anthropic, Google, Microsoft and OpenAI on the challenge, and will announce the winners at the Defcon hacker’s convention in 2025.

But don’t forget it’s an arms race. Bad actors can access AI, too.

## 4. Protect the Public's Rights

Despite its great potential for good, AI has revealed a potential to do harm.

Systems based on AI can reproduce societal biases and discrimination, for example in hiring, health care and credit decisions. Even more dangerously, some AI-driven facial recognition software used to identify criminal suspects has been unable to tell Black individuals apart, leading to wrongful arrests.

There are also serious privacy issues. Vast amounts of personal and other data fuel AI applications, making it easy to track and collect information about people's activities.


The key to avoiding these pitfalls is human involvement at all stages.

NIST counsels that people must be involved in all phases of a project, and policies must be in place to guard against bias and privacy violations.

To preserve the public's rights:

- Plan your project around program needs, not the technology.
- Include all stakeholders in your planning.
- Include diverse perspectives on both your business and technical teams.
- Practice good data governance.
- Assess the AI algorithms for potential privacy and bias impacts.
- Don't depend on vendors for security, privacy protection or elimination of bias.
- Evaluate and adjust AI programs on an ongoing basis.

The Biden administration's Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence underscores its commitment to leading by example in making AI safe: "The Federal Government should lead the way to global societal, economic, and technological progress. ... This leadership is not measured solely by the technological advancements our country makes. Effective leadership also means pioneering those systems and safeguards needed to deploy technology responsibly."



# How MLOps Helps Agencies Get the Most out of AI

An interview with John Dvorak, Chief Technology Officer for Red Hat North America Public Sector

Building a large AI data model from scratch is prohibitively expensive for most organizations. But the release of AI foundation models, particularly large language models (LLM), has allowed organizations to take advantage of broader community investments in AI.

But there's a catch: Agencies need to tweak and train these models to meet their specific objectives, maintain privacy and ensure regulatory compliance.

"These models are trained on a broad range of data, a wide understanding of language and concepts and images," said John Dvorak with Red Hat. "Fine-tuning is essential to maximizing the relevance of the model, the accuracy, the effectiveness for the specific use cases that you have."

## Understanding Foundation Models

Keep in mind that foundation models are purposely versatile and adaptable. Because they draw from a broad base of inputs, "they can be adapted for use across a wide variety of ranges and use cases," Dvorak said.

But that strength is also a potential weakness. In their raw form, broad-based foundation models are not ideally suited to support the nuanced needs of government agencies, Dvorak said, which deal in "unique libraries or vocabularies, terminologies and processes, which aren't necessarily captured or semantically linked in the model. They are also not tuned to address bias or handle concepts such as novel or emerging topics."

Those models also may not be adept at protecting sensitive or private data, or may not act in accordance with agency regulations and collection authorities.

## How MLOps Helps

By using their own data to adjust the foundation model, whether through fine-tuning or newer strategies such as Retrieval-Augmented Generation (RAG), agencies can drive more effective AI outputs. But any effort to maintain a model in production requires a secure, transparent and consistent process for making improvements over time.

This is where Machine Learning Operations (MLOps) comes into the picture. MLOps offers a streamlined approach to making iterative improvements to the model. "It's taking that model through its life cycle: from data collection, to training that data, putting it into production and then monitoring — then, going back and doing it again," Dvorak said.

Red Hat OpenShift AI provides a flexible and scalable platform for building AI-enabled applications. It includes all the elements of MLOps, empowering organizations to automate and simplify the iterative process of integrating machine learning models into software development processes, production rollout, monitoring, retraining and redeployment for continued accuracy.

With a flexible platform built on a scalable infrastructure, developers can hone the foundation models in support of AI applications that understand government's highly specific subject matter, and align to its particular operational requirements.

"OpenShift AI can run on prem, in the cloud, or on the edge of your network. The platform is plug-and-play: It's consistent, it's flexible, and provides all the components" to train, fine-tune, serve and monitor models, Dvorak said.



**Red Hat**



## Getting Comfortable With AI

To learn where agencies will take AI — and where AI will take them — GovLoop spoke with [Beth Simone Noveck](#), New Jersey's State Chief Artificial Intelligence Strategist and Director of [The Governance Lab](#) at Northeastern University.

"For 2024, I think we're going to see a lot more use of generative AI in government," Noveck said. "States are going to be more proactive in starting to issue policies of the kind [New Jersey Chief Technology Officer Christopher Rein [issued](#) last year] to embrace the use of responsible AI and guide their employees in how to use [it]."

States will clarify how existing privacy and nondiscrimination statutes apply to AI and consider new laws to protect the public's rights, Noveck predicted.

### We Will Learn to Use AI Responsibly

"We'll see a big push in 2024 to train public-sector workers in responsible uses of these tools that are careful and secure, to ensure that we're not violating anybody's privacy or putting up information that's inaccurate," Noveck said.

She pointed to New Jersey's new guidance on state use of AI. "First of all, the policy is very clear about not using personally identifiable information [with] these tools. Second, [it's about] ensuring that the use of these tools is disclosed. Third, when you ask ChatGPT to write something and are putting it on a website, [you need to review it] in the same way you would if you assigned an intern to help you write something," said Noveck.

The guidelines also call for ongoing training in the proper use of AI.

The Governance Lab offers [workshops](#) on AI in government on its [innovate.us](#) website, she noted.

### AI Technology Will Mature

Another trend is the increasing maturity of AI technologies, especially in the areas that matter most to governments. "I'm especially excited about the ability to ask generative AI to answer a question based only on a limited set of documents," Noveck said. "You can feed it with ... all the policy documents relating to service X or all the directions relating to how to apply it to service Y, and [it will] provide answers based only on that information."

"The ability to train on [a specific] set of documents helps to reduce a lot of problems we did worry about in the past and can help ensure greater accuracy and greater privacy," she added.

AI's growing multimodal capabilities also are making AI easier to use, Noveck said. Text-to-voice and voice-to-text applications mean AI tools can respond to spoken queries and produce spoken responses far more sophisticated than Alexa or Siri. "I have ChatGPT on my phone, and if I want to, I can go down the street asking it questions. It's pretty amazing," she said.

New Jersey has used a platform called Synthesia, an AI prompt-to-video generator, to make training videos. The AI is trained on a human actor's voice and motions to produce a credible video presentation.

"The other big change is just the integration of AI into the daily way we do things," Noveck said. It's moved way beyond autocorrect and autocomplete. For instance, Google Docs "can help you summarize your document or expand your document. Same with the Microsoft. It's becoming more standard under the hood," she said.

"It's one thing to go to ChatGPT and another to be just using your regular Outlook or Gmail and getting a lot of AI along with it," Noveck said.

# Meeting AI Where You Are

An interview with Ryan Macaleer, Vice President, Data and AI for the U.S. Federal Market at IBM

Traditional, task-specific AI has helped federal agencies improve operational efficiency, productivity and decision-making. With the emergence of generative AI, agencies are beginning to experiment with its ability to automate, augment and accelerate work. The White House [Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence \(AI\)](#) acknowledged the tremendous opportunity of AI and our shared commitment to harness it responsibly. Its focus on the federal government's use of AI systems is critical, as the government's massive purchasing power can set the bar for adopting trustworthy AI. But what's the best way for federal agencies to take their AI development and deployment to the next level with generative AI?

Most agencies will get the best value from adopting a use case driven approach, according to Ryan Macaleer of IBM.

## Adopt a Use Case Driven Approach

"Traditionally, AI projects were task specific and siloed. That technology enabled advancements in operational efficiency, productivity and decision making, but it required a lot of resources, computational power, and [skilled] data scientists to develop algorithms and rules for a specific task," he explained. That can be very expensive for one piece of work. "Each task had to have its own governance, its own data, its own structure," he said.

Federal agencies' experimentation with generative AI demonstrates that AI has moved beyond number-crunching and repetitive tasks. It's now capable of natural language processing

(NLP), grasping context and exhibiting elements of creativity. But one of the greatest challenges to adopting generative AI is knowing where to begin. According to Macaleer, developing an AI strategy means identifying how an agency can best use AI. In other words, identifying AI use cases.


IBM has identified three use cases that can offer agencies a quick return on investment: human resources, citizen services and application modernization. This approach allows agencies to develop uniform standards for governance, security, bias detection and other concerns. The AI is not confined to solving one particular problem, said Macaleer, "and you get the benefit of scale across an agency."

## Embrace Explainability and Build Trust

While the potential of generative AI is exciting, navigating the landscape requires a balancing act between progress and prudence. One common concern about generative AI is that it can seem like a "black box" whose workings are unclear. Knowing what data went into an output and why — explainability — builds trust.

Developing robust mechanisms to ensure the responsible use of generative AI technology is essential. "That is why we built powerful AI governance into [watsonx](#), our comprehensive AI and data platform, to give federal agencies the ability to manage the entire lifecycle of AI, including the training, tuning, deployment and ongoing governance," said Macaleer.

"It's an exciting time," Macaleer said. "I think we have so much more to gain than we do to be afraid of. So, let's embrace it together."



IBM

FOUR  
inc.

## Conclusion

This is the first of three AI guides planned for 2024.

In the upcoming guides, we'll focus on examples of AI in action at government agencies and ways to build your team's AI knowledge.

You can sign up for the rest of the series [here](#).



## 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](mailto:info@govloop.com).

## Thank You

Thank you to Elastic, IBM/Four Inc., Pure Storage, and Red Hat for their support of this valuable resource for public-sector professionals.

## Authors

**Lauren Walker**, Senior Staff Writer

**Candace Thorson**, Managing Editor

**Susan Kirby-Smith**, Senior Staff Writer

**John Monroe**, Director of Content

**Adam Stone**, Contributing Writer

## Designers

**Marc Tom**, Junior Graphic Designer

**Andrew Blake**, Motion Graphic Designer