

AI IN STATE GOVERNMENT

Balancing Innovation, Efficiency, and Risk



Katherine Barrett
Richard Greene



IBM Center for
The Business
of Government

AI in State Government: Balancing Innovation, Efficiency, and Risk


Katherine Barrett and Richard Greene
Executive Advisers, American Society for
Public Administration

OCTOBER 2025

Table of Contents

Foreword	5
Executive Summary	6
Introduction	7
Controlling Expectations	13
AI Advantages	17
Roadblocks	21
The Risks	25
The State of AI Governance	28
Observations from Experts	34
Conclusion: Vital Next Steps	35
About the Authors	36
Recent Reports from the IBM Center for The Business of Government	37

Excitement among state leaders for the potential of artificial intelligence is palpable; putting GenAI into production is still in an early stage, with conversations focused on the possibilities



Foreword

On behalf of the IBM Center for The Business of Government, we are pleased to present this new report, *AI in State Government: Balancing Innovation, Efficiency, and Risk*, by Katherine Barrett and Richard Greene of the American Society for Public Administration.

This timely report examines the rapid expansion of artificial intelligence (AI) applications across U.S. states, with a particular focus on the accelerating experimentation with generative AI (GenAI). State governments are increasingly exploring how GenAI can streamline operations, enhance service delivery, and support policy innovation—while safeguarding human judgment, transparency, and accountability that define public governance.

Through an in-depth review of current pilot projects, emerging use cases, and early implementation lessons, the authors offer a forward-looking perspective on how GenAI can serve as a collaborative partner for state employees. The report maps areas where AI can complement, augment, or automate tasks within diverse state functions, from public health and transportation to education and environmental management.

Key recommendations include fostering cross-agency learning networks, investing in targeted workforce training and upskilling, and adopting governance frameworks that balance innovation with responsible use. By following these strategies, states can cultivate a technologically empowered and resilient workforce in an era of rapid digital change.

This report extends the IBM Center's commitment to exploring the intersection of AI, workforce transformation, and government service delivery. It builds on prior Center publications such as *GenAI and the Future of Government Work*, which explores the transformative potential of GenAI in reshaping the workforce; *Navigating Generative AI in Government*, which outlines strategic pathways for integrating GenAI into public service; and *Digital Modernization for Government: An Implementation Framework*, which helps to create an evidence-based framework for digital modernization.

We are grateful to the authors for this insightful contribution, and believe the report will serve as a valuable resource for state leaders seeking to harness GenAI's potential while keeping the human element at the heart of public service.



Daniel J. Chenok
Executive Director
IBM Center for
The Business of Government
chenokd@us.ibm.com



Chris Shriver
Senior Partner-Vice
President, SLED Industry
Leader, IBM
cshriver@us.ibm.com



Executive Summary

"*AI in State Government: Balancing Innovation, Efficiency, and Risk*" delves into the rapid expansion of artificial intelligence (AI) applications across U.S. state governments, with a specific focus on the accelerating experimentation with generative AI (GenAI). The report highlights how state governments are exploring GenAI to streamline operations, enhance service delivery, and support policy innovation—while maintaining human judgment, transparency, and accountability.

Through an in-depth review of current pilot projects, emerging use cases, and implementation lessons, the report offers a forward-looking perspective on how GenAI can augment the work of state employees. AI can enhance, complement, or automate tasks within diverse state functions, from public health and transportation to education. The analysis emphasizes the importance of fostering cross-agency learning networks, investing in targeted workforce training, and adopting governance frameworks that balance innovation with responsible use.

The report reflects the importance of AI governance, recognizing that technology alone is not always the answer. Governance that includes transparency and public accountability can improve public trust. Various governance models are outlined, such as Ohio's structure requiring agencies to submit any generative AI use case to a central council for approval, and Texas's more decentralized approach.

The report also addresses the risks associated with GenAI, including hallucinations, biased data, and privacy and security concerns. The authors stress that while AI can augment workforce productivity, human judgment is necessary to address these and similar risks—like privacy issues associated with exposing sensitive data to public AI models.

Key recommendations include:

- **Invest in targeted workforce training:** Ensure that government employees at all levels, as well as the public, are adequately trained to understand the potential benefits and risks of AI.
- **Adopt governance frameworks that drive innovation and responsible use:** Implement evidence-based policies, use guidelines, and other measures to ensure effective and safe AI development and use.
- **Ensure transparency:** Preserve accountability and build trust by maintaining visibility into AI projects.

By following these strategies, states can cultivate a technologically empowered workforce that is resilient in an era of rapid digital change, while keeping the human element at the heart of public service.



Introduction

The use of artificial intelligence has greatly expanded in the states, with the number of experimental uses of Generative AI (GenAI) growing every day. “I think everybody is chasing AI one way or the other. It is a headliner for every conference,” said Alan Shark, co-chair of the National Academy of Public Administration’s Standing Panel on Technology Leadership and a nationally known expert on AI.

“This is so transformational. The impact of this is beyond what states have seen, and certainly what CIOs have seen in the in the past,” said Doug Robinson, executive director of the National Association of State CIOs (NASCIO).

Artificial intelligence is nothing particularly new, and the phrase AI is often used to cover all its variations. For some time, it has had the capacity to analyze or classify data to help governments perform a variety of tasks. In recent years, GenAI emerged with the ability to create new content (like text, images video and code) rather than simply processing information.

The next step, agentic AI, takes things a step further. Agentic AI uses machine learning models to move toward an objective—like the elimination of unproductive regulations—taking more autonomous actions, although human involvement is still important. The following use case outlines how Virginia used agentic AI to streamline government regulations.



VIRGINIA USE CASE

Streamlining Government

While a growing number of states are working on streamlining operations to use tax dollars more efficiently, Virginia is the first to announce it is utilizing agentic AI to promote efficiency in the regulatory space. On July 11, 2025, Virginia's Governor Glenn Youngkin issued an [executive order](#) launching a pilot program for a first-of-its-kind agentic artificial intelligence regulatory reduction tool, which will help reduce regulatory burdens and keep regulations and guidance documents both streamlined and up to date.

According to [IBM](#), "Agentic AI is an artificial intelligence system that can accomplish a specific goal with limited supervision. It consists of AI agents—machine learning models that mimic human decision-making to solve problems in real time." In a system using multiple AI agents, each performs a specific task to reach the goal, with "AI orchestration" coordinating these efforts.

This new tool has a variety of potential uses, including doing cost benefit analyses on state regulations. Another potential utility allows the virtual agent to break a statute down into its individual pieces to determine if an individual regulation is actually mandated by the statute.

As Reeve Bull, director of Virginia's Office of Regulatory Management, explained: If the statute says there will be a \$100 fee for a service and the regulation calls for that, "there'd be a 100 percent match. But if the statute says to charge a \$100 fee and the agency is charging a \$200 fee, then they're violating the statute."

The agentic AI system in the commonwealth enables this kind of independent analysis, point by point, going through hundreds of thousands of separate provisions within the overall code, and making determinations in each case.

Then, of course, said Bull, "there always has to be a human in the loop. We've been explicit on that, from the outset, that the AI is merely a really robust research assistant that's doing, in a matter of minutes, what it would take a human being hundreds of hours to do. But it's not, itself, changing regulations. It's just giving recommendations for possible changes."

The agentic AI initiative in Virginia has the capacity to go still further in streamlining than just making certain that statutes line up with regulations. It can also compare Virginia to other states to see if its regulations are out of line with what's done elsewhere.

"Virginia had long required 1,500 hours to become a cosmetologist, which was then cut to 1,000," Bull said. "This prompted the Governor's office to consider whether there were similar opportunities for reductions in that or other professions. To see if that's the case, 'the AI tool can go in and ask, 'Okay, what do the surrounding states require? What does Tennessee require? What does West Virginia require? What does Kentucky require?' And then it can quickly produce a report and say here's where the surrounding states are, and you're either low or high compared to those.'"

Though the excitement among state leaders for the potential of artificial intelligence is palpable, putting GenAI into production is still in an early stage, with conversations focused on the possibilities and many officials in state agencies still scratching their heads because they don't have the vision, time, staff, money or governance systems ready to handle the risks or the complexity of organizational change. "I think this discussion will be very different a year from now," said the National Academy's Alan Shark.

Precise data about the number of states at various stages of progress are hard to come by. When asked to supply data about where his state stands in developing prototypes, pilots and full-scale production of AI uses, one leader of a state's AI program said, "I don't want to give you a number, because it would be out of date by the time you print."

Generally, in early September, NASCIO's Robinson characterized about a third of states as being "the leaders while a third are in the middle and a third are lagging. They don't have enterprise policies. They don't have a roadmap."

States are in a transition period from proofs of concept to pilots and for a small portion actual implementation right now, although some remarkable projects are already in or near realization, "I see calendar year 2026 as the major year of a lot of implementations, but we're seeing that start now," Robinson said.

Utah, for example, is in that top third. "We are leaning in very aggressively. As a CIO, my role is about using AI in state government to do its job better," said Alan Fuller, Utah's CIO. "I want to give our employees superpowers. We can train generative AI models to do things that can be really beneficial to employees and help them be more productive in their jobs." The following use case shows how GenAI will be making life easier for the state's employees and taxpayers.

UTAH USE CASE

Help Come Tax Time

The Utah Tax Commission is responsible for collecting, processing, and enforcing all the state's taxes. As is the case with other states, this office is besieged with queries from residents who do not want to make a mistake on their filings. To deal with this, the state currently employs 200 call-center agents, with knowledge based on experience and training.

Since there's a fair amount of turnover in these ranks, training is an ongoing process, and even then, the newer agents tend to be less prepared to respond quickly and accurately.

So, the state ran a pilot utilizing four vendors who used a large language learning model to produce an AI tool that could tap answers to the 366 typical questions that people ask when they call in, based on an initial software script analysis. Those answers were then analyzed by knowledgeable tax and subject matter experts to judge the result compared to what a knowledgeable call center agent would say.

The result: 92 percent of the time, at least one of the models answered the question as accurately as a knowledgeable agent. Vendors thought they could get closer to 99 percent by fine tuning.

The state picked one vendor to roll out the use of the AI tool in pilot form on a small number of agents, and the model is being further refined. “And we expect to roll it out to all of the agents to give every agent a very powerful digital assistant that can help them answer questions more accurately and faster,” said Utah’s CIO Alan Fuller.

The creation of an AI chatbot that could be accessed directly by the public is still a work in progress. “We’re aggressively working to get a chatbot put together,” said Fuller, “but we’re going to keep it internal until we get to the point where the tax commission feels confident that it is giving accurate enough and consistent enough answers that they can make it available directly to the public. If you come to the tax commission for help with your taxes, you need to be able to trust that you’re getting good information back.”

Currently, chatbots dominate the AI scene, but in state governments, they are mostly being piloted for use behind the scenes, helping staff to find the answer to questions rather than connecting the chatbots with callers themselves. Arizona is piloting a GenAI chatbot to assist field workers for the Department of Child Welfare in getting quick research on state and federal policies that are relevant to a case. The pilot has been helpful in demonstrating to field workers that AI could make their work lives easier.

The numbers of innovative and intriguing uses are multiplying. Washington state, also a leader, is piloting practical AI applications beyond chatbots, including computer vision for environmental monitoring and AI-enabled automation to speed routine workflows.

In environmental work, for example, computer vision in Washington now helps pre-screen imagery for tasks like estimating fish passage. Field experts remain in the loop, validating outputs and making decisions, and the state is investing in staff training so people know how to evaluate these tools and their results. “The goal is faster, more consistent monitoring so scientists can spend more time on analysis and on-the-ground action,” said James Galvin, who leads Washington’s AI and Emerging Technology program.

Agencies are moving at different speeds when it comes to using AI. As he explained, “Some agencies are more in a learning stage while others are actively evaluating and testing AI’s potential.” That includes augmenting staff and improving service, productivity and safety—for example, using AI-assisted wildfire detection. The following use case shows how Washington’s AI Vision cameras are helping to locate the beginning of wildfires.



WASHINGTON USE CASE

Wildfires

There are a great many concerns around artificial intelligence voiced by the public and state staffers, but few uses of the technology have gotten more immediate buy in than Washington state's use of AI to control wildfires and predict where wildfires are likely going to occur.

The Washington State Department of Natural Resources (DNR), in collaboration with various agencies including Washington Technology Solutions (WaTech), has implemented an AI-powered wildfire detection program to modernize fire situational awareness and response. It's already operational and has gone from pilot to production over the last 18 months.

According to James Galvin, the lead for Washington's AI and Emerging Tech Program: "States like Washington haven't used rangers with binoculars for decades. Today, wildfire detection helps fill that gap with 30 high powered cameras on tall towers looking across the state, helping human operators in identifying potential wildfires. They're using AI in these tools to more accurately identify what is genuinely wildfire smoke versus dust from a dust storm or low hanging clouds. They're able to continually improve the algorithms. So, using AI, they're able to really improve the quick identification of wildfires and have less false positives."

With that information in hand, using a dashboard that coordinates the information by longitude and latitude, the state can deploy the personnel to extinguish the flames before they have the chance to spread. "It's really about getting to those wildfires quicker, which ultimately reduces the costs of fighting the fire in manpower, resources, aircraft, as well as potential damage," said Galvin.

In 2025, five more stations are planned to go operational, with plans to provide new features including additional AI-driven intelligence layers. The state is building intelligence layers for jurisdictional boundaries, utility lines, historical sites, and more to further improve multiagency use of the data.

The state started piloting this effort in 2023 and has increasingly added more AI analytics and map overlays, as well as increased the number of cameras and looked at not just the wildfires that's in the forest, but also urban wildfires, like those that savaged the Los Angeles area at the beginning of 2025. That fire began with just a spark. The state is now putting up additional AI vision cameras on utility poles owned by Puget Sound Energy. Washington's Galvin said, "If there is a spark, we want to be able to make sure that the information is shared and we're able to respond."



In the pages that follow this report will provide readers with a thorough sense of a number of the most important topics to leaders in all corners of state government, including people in agencies making advances in GenAI; legislators who need to understand the power and the risks of these tools; chief information officers and their staffs; government employees whose work is increasingly becoming enmeshed with AI; and ultimately the public.

Topics covered will include:

- The benefits of moving ahead with new AI tools
- Obstacles in the way of more AI development
- The risks of generative artificial intelligence
- Current state AI governance
- Ongoing issues with new policies and approaches
- Observations about early use of new AI tools

Additionally, this paper features a series of sidebar/boxes which delve deeply into some of the most intriguing projects that states have embarked upon with generative artificial intelligence.



Controlling Expectations

For most states, caution has been key. “We want to move fast, but we don’t want to break things,” said Daniel Egan, director of Communications at the Pennsylvania Office of Administration.

His sentiments were mirrored by officials from a number of other states, concerned, not just about the potential of AI but the critical policy considerations that accompany that progress. “People have been respectful of trying to understand what those things are before they rush to solutions,” said Mark Raymond, CIO in Connecticut.

Of course, the power of AI is such that there’s always a risk that, if its uses aren’t carefully vetted as are the data that underlies its conclusions, it can create unfortunate outcomes.

Jason Snyder, Massachusetts CIO and secretary of the Executive Office of Technology Services and Security, neatly summed this up. “You cannot count on AI to enhance a flawed process,” he said. “If your process and data are not up to standard, AI will only amplify these shortcomings.”

For example:

- A benefits application process already riddled with bias will only deny people faster if AI is layered on top.
- A procurement process with inconsistent rules will just generate inconsistent outputs more efficiently.
- Inaccurate or incomplete data sets will be enshrined into models that scale bad decisions.

And the stakes are rising with AI agents entering the market. These are systems capable of acting with autonomy, triggering workflows, and making consequential recommendations. Without guardrails, they risk running ahead of human oversight. If a chatbot gives bad information, the harm is limited; if an AI agent autonomously denies benefits, approves a payment, or escalates an enforcement action, the impact is systemic and immediate.

Governance Is Key

AI governance requires humility: admitting that technology alone is not the solution. Without that stance—and without the literacy to question AI's role—governments risk building faster, shinier versions of the very systems that people already distrust.

True AI accountability means ensuring that models, data sources, and decision protocols remain subject to public scrutiny. Without the kind of transparency that the Freedom of Information Act provides for documents, “responsible AI” risks that the general public will see AI as a threat. This powerful new technology can become a black box, which affects everyone's lives without any understanding of what it is and how it works, and that absence of understanding can lead to the erosion of government trust.

The states are influenced by a wide variety of organizational, management, political, and policy factors. “AI is not a standalone in a vacuum,” said Vanitha Zacharias, who leads Ohio's AI and AI governance. These efforts connect with executive leadership and all the different governance areas; data security, privacy, legal.

“It's not like we put in AI and that's all we concentrated on. It's always a holistic approach, where AI components are one part of the story,” said John Harrison, director of Information Technology at the New Jersey Department of Community Affairs. The following use case describes the Department of Community Affairs “Ask Claudia” AI assistant.

NEW JERSEY USE CASE

“Ask Claudia”

For decades, human resource officials have been worrying about the difficulties of implementing effective knowledge transfer practices leaving agencies with a chronic malady: institutional brain-drain.

New Jersey, like many other states, is suffering through a major loss of intuitional knowledge as long-tenured employees retire from public service. “There are people retiring with 40 years of experience and the next person has 10 to 15 years at most,” said John Harrison, director of Information Technology at the New Jersey Department of Community Affairs (DCA), a state agency that offers programs that serve local governments, businesses, and residents in New Jersey.

New Jersey is addressing this governmental deficit with its “Ask Claudia” AI assistant, which is currently in development and is expected to be released in October. The AI assistant has been named after a high-level supervisor at the Department of Community affairs, named Claudia, who has developed an unparalleled depth of knowledge and experience in the science of judging whether applications qualify for benefits by applying program policies and regulations to individual cases.

Currently, it's a common occurrence for other employees to tap Claudia on the shoulder to ask questions that she can answer off the top of her head but might take hours of research to uncover otherwise.

If Claudia wasn't available to them, what would staffers do?

"We're creating the 'Ask Claudia' AI assistant," explained Harrison, "so when Claudia retires in the imminent future, people can still ask Claudia, but they're asking the AI assistant." Naming the tool after a real living human being, with whom employees are already comfortable interacting, may well help them make the shift more comfortable in using AI for similar purposes.

"We'll try to move as much of her brain into the system as we can," he continued. "And when she does leave, we're going to have a system that can continue to grow with knowledge over time with more data points as more cases are processed. As more junior staff 'ask Claudia' a question, that will increase the AI assistant's knowledge." This pilot may lead to the same approach taken to capture institutional knowledge held by senior retiring employees in other areas.

"We want to build institutional knowledge in the AI, so we don't have that institutional knowledge drain," said Harrison. "We want to capture it. We want to put some boundaries around it, build a little dam and keep all that information in the institution."

This AI assistant is intended to be used only by government employees and is not intended to be accessed by the public in the near future. "We're trying to become experts on AI before we worry about the security and privacy issues that come with using it to interact with the public," Harrison said. "The adjudicators on benefit applications will be able to use this internal AI assistant to ask policy questions. It has access to our policy documents and previous case data."

Harrison added, "We're keeping humans in the loop, and AI isn't making any decisions or adjudicating any benefit case decisions. It's just providing policy information and relevant information for the case being processed."

As author and consultant Peter Drucker said decades ago, "In today's environment, the most important resource is knowledge." Perhaps systems like 'Ask Claudia' will be the equivalent of placing that resource in the bank and watching it gather interest.



As AI advances, it's clear that the questions of governance are of paramount importance. Who, for example, is ultimately responsible for making the decisions about whether agencies should be allowed to proceed with grand new visions for use of AI? Texas, for example, leaves open a great deal of latitude to its agencies (not unexpected in this deeply decentralized state), while Ohio has a rigorous organizational structure that requires agencies to submit any generative AI use case to a central council for approval.

Said NASCIO's Robinson, "The hard part is governing it and funding it, and ethics and responsible use and legal, privacy and security issues. That's the hard part. It's easy to buy it and deploy but it's hard to make sure you stop bad things from happening when you're doing that."



AI Advantages

Competitive benefits

Generative artificial intelligence brings multiple advantages with an array of benefits starting to materialize in states that use it wisely and well. The list of benefits that AI advances can provide just begins with:

- Budgetary solutions as federal funding declines
- Translating complex documents into simple English
- Aiding government staff in understanding the needs of the public
- Helping the public access information and services
- Providing faster call responses for taxpayers
- Streamlining work processes
- Improving agency communication and collaboration

With these benefits in mind, AI is not just an appealing advance to technocrats but to agency heads, other government leaders, and government employees who see tedious routine work diminishing and more time available for helping the beneficiaries of services, strategic thinking, and working to solve countless enduring state challenges.

Gen AI is affecting government services in several ways, described below.

Solving resource issues

“AI is not only something that can be looked upon as an efficiency tool, but it can be looked upon as a way of possibly saving money,” said the National Academy’s Shark. “One of the most expensive parts of government is labor. Up to this point, most people have argued somewhat persuasively that the movement towards AI adoption is not about replacing humans, it’s about augmenting them. And I’m comfortable with that. But I think we’re going to see less entry level positions being filled, and when people leave, there’s going to be more pressure not to fill those positions.”

While the specter of future job loss has alarmed many observers, there are still a number of government services that are understaffed; state performance audits often cite staffing shortages as a reason that policy goals are unmet or that basic processes take too much time.

For example, procurement is still a field that continues to experience shortages of experienced staff and a surfeit of compliance-oriented work that could be automated. “Technology is a way to impact and address talent shortages,” said Rick Grimm, chief executive officer at NIGP: The Institute for Public Procurement, “because procurement, by its very nature, is very process driven—from compliance and bid evaluations to the solicitation process.”

Driving efficiency and effectiveness

The Department of Community Affairs in New Jersey has been piloting the use of AI in the area of utility benefit applications. “We’ve dramatically reduced our time to review an application,” said director John Harrison. In that department, AI can now process 2.5 million documents for about \$40,000 a year. And the adjudication time showed an 88 percent improvement.¹

In Arizona, the Department of Child Welfare was spun off of the Department of Economic Security about 15 years ago and faced the common problem of dealing with data that was spread over a number of entities. This required a fair amount of time from individuals tasked with pulling data together in a way that was useful.

“Communication and collaboration were very manual and cumbersome,” recalled Frank Sweeney, CIO of the Department of Child Safety there. That was a clear driver in encouraging the department to build an AI program that cuts the hours field workers spend on administrative work and provides them more ability to work with families and children.² “We know there’s a direct correlation between having more time with families and kids and a positive outcome,” he said.

Translating technical verbiage into plain language

In a Pennsylvania pilot, 175 employees were given ChatGPT Enterprise over the course of a year. The following use case shows what the state learned about AI from its employees.

1. “Delivering help where it’s needed, when it’s needed,” case study, New Jersey Department of Community Affairs and IBM Consulting. <https://www.ibm.com/case-studies/new-jersey-department-of-community-affairs-consulting>.
2. “Arizona Department of Child Safety Simplifies Processes with GenAI,” case study. <https://www.ibm.com/case-studies/arizona-dcs-consulting>.

PENNSYLVANIA USE CASE

Learning From Employee Experience

In January 2024, Pennsylvania launched a pilot program to help explore the potential of AI while gathering valuable data from employees about their experiences. This move forward was a natural extension of a [September 2023 executive order from Governor Josh Shapiro](#), which outlined his hopes and concerns about the future of AI in that state. This [first-of-its kind GenAI pilot](#) was created in collaboration with OpenAI and Carnegie Mellon University's Block Center for Technology and Society.

The pilot involved 175 employees from 14 different agencies, including a wide range of workers in different disciplines. They had access to ChatGPT Enterprise (a more sophisticated version of ChatGPT, for which there is a licensing fee), for about 12 months. Within a week of the launch, the state's Office of Administration put up a new website to help guide state employees as they took this journey.

Participating employees were able to utilize this tool for a variety of tasks including writing emails and job descriptions, reducing longer documents into more succinct ones, coding and debugging, research, and streamlining hiring processes.

It was “really getting to work with our employees directly to understand the types of uses that seemed readily apparent to them as they had exposure to AI tools,” said Harrison MacRae, director of Emerging Technologies for Pennsylvania. “That’s been a really valuable learning for us throughout this project.”

The findings were very positive, and contributed to the state's enthusiasm for the future of AI in its borders. “This was a pilot with a limited number of licenses, and the state has continued to scale access of AI tools to employees,” explained MacRae.

For example, compared to other new tools and technologies “our employees really enjoyed using these tools,” he said. Although some started skeptically about this enterprise, “about 85 percent had a positive experience using ChatGPT in their work.”

The average employee reported saving about 95 minutes per day that they used ChatGPT. This dramatic finding was key in giving state leaders an indication of the tool's value.

One of the lessons learned from the pilot was the need to work with employees who begin with a wide variety of proficiency levels, and who will be using GenAI in different ways. This, in turn, led the state to partner with Innovate U.S. for ongoing training.

As Governor Shapiro was quoted in a [March 21 press release](#): “Pennsylvania is leading the way in responsibly integrating AI into government by giving Commonwealth employees access to tools that enhance efficiency while ensuring people remain at the center of decision-making. This pilot program showed that when used thoughtfully, generative AI can help employees save time, streamline processes, and improve services for Pennsylvanians. But let me be clear—AI will never replace our workers. Instead, we’re equipping them with the best tools to do what they do best: get stuff done for Pennsylvanians.”

Observers noted that the state benefitted through better communication with employees who learned about ways the new tools could turn different types of documentation, technical specs, and legal jargon into plain language that was relevant for an end user audience.

“That was a use case that we saw constantly in different teams,” said Harrison MacRae, director of Emerging Technologies there. “Plain language is very simple on its face but is also essential to a lot of government forms and, work processes.” In effect, GenAI helps to make different types of services or forms easier to understand and navigate both for government workers and for communicating with the public.

In Connecticut, the state went live in mid-summer with pilot chatbots in use for Medicaid and the Supplemental Nutrition Assistance Program (SNAP). The chatbots can trace application history and shorten the wait time for talking to someone or to find out the status of an application. “They can go and check online and be able to understand the question and connect to the underlying system,” said the state’s CIO Mark Raymond.

Addressing security issues

Utah is using AI to combat cyberthreats. Explained CIO Fuller, “We have a vendor provided product where we take all of our network traffic and terabytes per day of information. We dump it into this tool, and we have trained AI agents that go through it and identify any snippets of code on the network that may appear to be malicious. Then they kick out a service request that a human can look at. That has been very effective for us in finding risks in our network where we have malware or malicious software that somebody has managed to get into our network. We can identify it quickly. So, it’s kind of the digital equivalent of finding smoke where there’s a forest fire.”

Communication and collaboration

A beneficial side effect of working on GenAI solutions was that it has prompted more collaboration between agencies and central governments to ensure responsible use while working on improving services.

“Typically, in a state as large as California, we have a tendency to go off and work on our own problems, individually, and have our own missions,” said California CIO Liana Bailey-Crimmins. But working on GenAI “was so refreshing that it brought all of us together in a laser-focused manner to achieve outcomes for the benefit of Californians. I mean, generative AI is great, but I feel like it was really the relationships, the outcomes, the ability to do things the way we’ve never done it before, which is what excited me.”

“In a state as diverse and dynamic as California, it’s easy for teams to focus on their own missions. We made a conscious decision to approach GenAI in a united and focused way with an opportunity to deliver for Californians,” said Bailey-Crimmins. “What excites me are the relationships and the outcomes we built together. With that kind of collaboration, we’ll stay focused on what matters—delivering values to our residents.”

Bailey-Crimmins and Jonathan Porat, the state’s chief technology officer, describe the effect GenAI has had on increasing interdepartmental sharing and suggest that it may contribute to fewer government silos.



Roadblocks

While the many potential benefits of GenAI are intriguing, the path to implementation is also full of challenges that need to be dodged in order to move forward safely, productively, and responsibly.

While these obstacles mostly slow down progress, in some cases they may mar a state's competitive advantages. "In an era of enhanced technology options, obstacles on the road to GenAI hamper opportunities to gain greater efficiencies," said CEO Rick Grimm at NIGP. "I've heard a couple of instances where governments are not allowing their procurement teams to use AI for any reason. At the same time, they are underfunding resources to hire staff needed to perform procurement tasks. Limiting hiring practices while limiting technology advances is really short-sighted."

Here are some of the challenges that states face as they work out how to take advantage of the benefits that AI offers.

Poor quality data

AI has sparked a push to improve data quality in Arizona. While the state's data has generally been good, quality was still a concern "because bad data gives us bad outcomes or bad decisions," said Arizona's Frank Sweeney.

Poor data can easily stymie the development of state use cases. In Massachusetts, a Center of Excellence has been set up to evaluate and mitigate risk in AI use cases that come from agencies. "We have found some of the use cases have ultimately gone back to step one because the data sets were unclear, so they're generating inaccurate information," said CIO Jason Snyder. When that's the case, the answer that comes to the agency is: "You can't go live. Go clean your data and when you get that done, then talk about AI," he added. The following use case describes a Massachusetts GenAI that helps parents search for a preschool.

MASSACHUSETTS USE CASE

Searching for a Preschool

In this exploration of the use of AI in the 50 states, there was no single utility that was more common than the use of GenAI chatbots, which help to access caller questions with accuracy and speed.

Currently, many chatbots are internal facing and are used by state staffers in order to help them easily answer questions that come from the public. Few have yet used sophisticated GenAI chatbots directly for the public's use, but a pilot from the Commonwealth of Massachusetts shows how it's accomplishing this with a chatbot that provides information on preschool openings.

This idea was sparked by a state contest in which agencies were encouraged to develop use cases. "I said, work with your teams and see what you come up with," said Jason Snyder, the Commonwealth's CIO. "The winning one was a chatbot that would give parents the ability to search for a preschool that meets their needs, including, for example, a preschool that could serve a special needs child."

When the tool comes into public operation, it will do the preschool searching for the user. It will provide information based on what a resident asks for and come back with preschools that have openings. "It meets the needs you defined in your natural language request and provides information on how to contact the school," said Snyder.

The creation of the chatbot is possible because Massachusetts tracks data on pre-school openings, with individual schools sending data to the state. "This was an organic bottom up (idea), and the team came up with it. I was like, 'I love this. This is brilliant,'" said Snyder, who has five sons and knows how much time parents spend trying to find the right school for their children. "I wasn't involved in it until they presented it back to me at the contest."



A complex vendor environment

With proliferating competing products and aggressive marketing to win government clients, states can find themselves perplexed about what to do in order to purchase exactly what they need. In Massachusetts, Snyder described the difficult times that states were having wending their way through the white noise of sales pitches that promise “We can do this.” Or “We can do this better.”

A additional channel for the Commonwealth of Massachusetts, given its educational resources, has been to turn to higher education institutions that have existing programs and are willing to partner. This option, Snyder said, also offers a lower cost alternative. An example he offered is the Worcester Polytechnic Institute, which is working on improving the state’s data set, a crucial task that’s needed for ensuring that a GenAI search for answers does not pick up false information.

Development and maintenance expense

While artificial intelligence use cases have sparked widespread hopes of efficiency savings, the development and maintenance of AI systems can be expensive and the states currently lack sustainable funding, according to NASCIO’s Robinson. That’s something that’s needed “if this is going to be transformational, save dollars, and streamline services,” he said.

“One of my fears,” said Shark, of the National Academy, “is that the cost of AI is going to go way high. It has to, with all the billions of dollars that have been invested in it.”

Sometimes, the idea for an AI use case gets dropped because of the expense of maintaining it. For example, when an agency in Ohio was considering a PDF summarization tool, it decided against going forward because there would be a monthly per user cost that Vanitha Zacharias—who leads the state’s AI and AI governance efforts and heads the Investment and Governance Division within the Office of Information and Technology—said was \$4 per month per user. For the agency, that meant the AI tool would not yield the right return on investment.

Worker inexperience

A widely recognized concern is “the readiness of the state workforce” to responsibly and capably deal with unfamiliar AI tools and a potentially dramatic change in work processes. While 71 percent of state CIOs told NASCIO that they are training workers to rectify skill gaps, there are still questions as to how deep the training goes.³ AI training needs compete with accelerated training demands in many areas. Training can also be expensive and is often the first to be cut when there are signs of budget trouble.

So, while there is a lot of training effort going on, there is also a clear skills gap. “We do need to train our staff up, skill them in AI, what to look for, (so that the) human in the loop validates outputs and responses from AI,” said Washington’s Galvin.

Ohio has partnered with Innovate U.S. to roll out responsible AI training for public professionals with seven modules that take employees through different issues that cover how to interact with AI, as well as how to protect data and not risk putting it out publicly. It provides the basics of how to interface with AI intelligently.

3. Full data from the 2025 NASCIO CIO survey will be available on the NASCIO website on tk.



The state has also created “training environments” for the workforce for hands-on training in-house, and then agencies have their own training specific to use cases rolled out to help employees securely adapt to generative AI interfaces.

Of course it’s not just the technical staff who need training. Legislators, frontline workers, and even the public must be equipped with the capacity to understand both the potential and the risks of AI. Without this broad understanding, oversight will fall exclusively in the province of technocrats and though they play an important role, the enormous potential of AI to benefit society means that a number of players must be literate in its use.

When staff isn’t well educated about AI—which increasingly is being used by them on a day-to-day basis—there’s a real risk that they’ll misuse this tool, potentially jeopardizing privacy and the security of government operations.

Federal encroachment on state decision-making

Several months ago, as part of the upcoming budget bill, Congress pushed for the inclusion of a stipulation that would have severely limited the states’ ability to regulate AI.

Although that notion was not included in the bill when it was passed in early July,⁴ the threat lives on. As Robinson pointed out, conversation is still taking place in Congress about introducing separate AI moratorium language to limit state regulation. In addition, the Trump Administration issued an AI Action Plan in July 2025. Its message: “The federal government should not allow AI related federal funding to be directed towards states with burdensome AI regulations that waste these funds, but should also not interfere with the state’s rights to pass prudent laws that are not unduly restricted to innovation.”⁵

The assignment of determining how to define “burdensome” was given to the FCC.

Even absent any clear directives, states fear that if they put in regulations to govern the use of AI, and the federal government objects, it could endanger the flow of federal money.

4. P.L. 119-21, An Act to provide for reconciliation pursuant to title II of H. Con. Res. 14, enacted by the 119th U.S. Congress; signed by President Trump on July 4, 2025.

5. “Winning the Race: AMERICA’S AI ACTION PLAN”, The White House, July 2025. <https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf>.



The Risks

GenAI is like a brand-new medicine for state governments, with the potential to help cure a great many budgetary, policy, and management ills that assail them. But like any new medical treatment, it's important to be aware of the unintended side effects before plunging in.

One primary concern goes under the evocative name “hallucination.” According to a [September 2023 IBM Report about hallucinations](#), “Generally, if a user makes a request of a generative AI tool, they desire an output that appropriately addresses the prompt (that is, a correct answer to a question). However, sometimes AI algorithms produce outputs that are not based on training data . . . or do not follow any identifiable pattern. In other words, it ‘hallucinates’ the response.”

As time goes on and the technology improves—it's likely that there will be fewer inaccuracies, and in the meantime, all the state CIOs interviewed for this report emphasized the need to “keep the human in the loop,” to validate information that AI produces.

What's going to happen to the workforce?

Since artificial intelligence can do repetitive tasks far more quickly than human beings, a widespread worry focuses on the effect its time-saving tools will have on state employee jobs.

As Utah CIO Fuller said, “There's a big question out there and a big question in the press about (whether) this will mean fewer employees needed? And maybe yes; maybe no. I think it's early to tell. (But) you have some technology CEOs out there on record saying they don't think we'll need as many employees to do the same job.”

As the shift to more sophisticated chatbots becomes routine, repetitive, boring calls are becoming the province of AI—with the idea of saving “really good people for the more complex things that they could solve, which would make their lives more interesting,” said Fuller. His goal in Utah is not to lay anyone off, but to “free up people to do other higher value things.”

One solution often recommended involves retraining displaced workers in other areas. But there are doubts. “That is a goal,” said Alan Shark, “but I don't think it's a workable strategy. It's aspirational, but I don't see it occurring operationally. You cannot just take everybody and say, ‘We're going to retrain you.’ We have to take a couple to retrain that are the better ones and say to the others. ‘Sorry. There's no room for you here because we can't afford you. Nothing personal.’”

At least for the moment, the more extreme fears about the loss of jobs hasn't been evidenced. "AI is never going to be a substitute for human judgement, compassion, empathy, and creativity that our workforce brings," said Massachusetts's CIO Snyder. "AI, at its most successful is a tool that augments the workers for the Commonwealth, right? We are seeing the benefits of AI as a safety net, but we are not seeing any evidence that AI can replace the person."

But Fuller points out that legislators may not see things the same way. "There's going to be pressure on budgets," he said. "I don't think those fears are unfounded in many cases. Our push is to generate efficiency so that we can reinvest in higher value things. But ultimately, will the legislature come to us and say, 'Yes. But we're cutting your budget 5 percent . . .'"

"You see what we're grappling with," he continued. "We want to do tremendous things with AI and be more efficient and drive productivity. Some of that productivity may come in the form of cost savings by a reduction of staff, enforced by the legislature, despite my focus and best efforts to make it a way that we can free up funds to invest in higher value return on investment projects. . . . Nevertheless, driving toward greater efficiency is a good thing."

Recently the IBM Center published a [study on the effects of GenAI](#) on government work, led by William Resh at Georgia State University; this report found that Gen AI has the potential to complement and improve work for a far larger number of government jobs than to replace them.

Built in biases

While no experts indicated that governments are purposefully building biases into the algorithms that underly AI, unintended biases can easily creep in. As one member of a roundtable assembled by the IBM Center for The Business of Government late last year commented: "AI requires that we develop algorithms, and if those algorithms have implicit bias built into them, it's going to exacerbate, not improve, the situation."

Ultimately, the problems with bias in AI become particularly acute when "the decision-making process is influenced inappropriately by the use of AI," said another member of that roundtable. And when bias creeps into the systems the burgeoning use of AI can lead to harm. As another participant pointed out, "Trust in institutions is dwindling," and when biases are uncovered, that phenomenon will only be exacerbated.

Privacy and security concerns

"If you include sensitive information and data while using ChatGPT, that information is now public," said Snyder. He underscored the critical risk of exposing sensitive data to public AI models, which may inadvertently lead to privacy breaches. When data is shared with these models, it can be stored and even used for further training, raising significant concerns about data security and confidentiality. Snyder emphasized the importance of being cautious and ensuring that sensitive information is not transmitted to such platforms without appropriate safeguards.

A number of states have banned the use of AI transcription services by state employees. Texas CIO Amanda Crawford said, "When we get into meetings with vendors . . . we request that those get turned off," adding that most state agencies don't use AI transcription for recording their own meetings either. In addition to concerns about state information becoming part of GenAI training, using AI transcription could violate state records retention laws.

It's not just AI projects developed by the state that raise concerns, but also those that are primarily in the hands of vendors or even nonprofits with whom the states collaborate. Working with third parties can make states vulnerable to external organizations that are dealing with state data.

Vulnerability is also increased based on an awareness that bad actors can use GenAI tools, too. Phishing attacks have become more sophisticated when they are supported by or enabled by GenAI.⁶ For example, foreign adversaries have been able to use GenAI to capture information about a state agency and use it in phishing attacks aimed at state employees.

In a survey of the state-run Texas AI 800-plus person user group (including agencies, cities, counties, school districts, health districts, etc.), avoiding privacy infringement was one of the topics most often requested by members—with questions like: How to ensure that there isn't leakage of state data? How do you prepare for leaks? What are best practices for shielding personal information from AI?

The user group has had workshops to discuss how to operate AI safely, how to deal with personal identifying information or health information in data, and how to look out for bias. The theme of these workshop sessions has been "Let us help equip you to establish governance for safe and secure and reliable AI implementation," said David Tucker, director of Strategic Digital Services who started the user group.

California has been particularly cautious about protecting the data that AI utilizes in artificial intelligence projects. For example, California CIO Bailey-Crimmins said that in dealing with two health projects, the state realized that even in publicly available data, a bit of personal information could sneak in.

With that in mind, California took a very cautious approach in the testing phase of one of the health projects mentioned by the CIO—the inspection of hospitals. "We had zero tolerance for risk," Bailey-Crimmins said, explaining that while the data was publicly available, it had been collected through hand-written notes by inspectors.

There was never a danger of a breach. Since inspectors might notice something that could expose even a small amount of personal information, the state chose to use dummy data during the testing phase rather than real data. "We know that once you lose someone's trust, especially at launch, you rarely get it back," Bailey-Crimmins explained. The state intends to build in enough safeguards so that real data can be used safely when this AI use case moves into production.

6. "AI-Assisted Cyberattacks and Scams," NYU, <https://www.nyu.edu/life/information-technology/safe-computing/checklists-guides/cyberthreats/ai-assisted-cyberattacks.html>.



The State of AI Governance

Without solid governance structures, governments run the risk of letting employees, or even departments, “do their own thing,” said the National Academy’s Alan Shark—a notion that is perceived by many to be a recipe for dramatic missteps on a potential path to success.

The need for strong policies, ethical guidelines, and guardrails for generative AI development was clear as soon as reports of GenAI power began to circulate widely in late 2022 and early 2023.⁷ Multiple states set up advisory boards and central organizations to vet use cases, work with agencies, and in some cases provide centralized state approval of what agencies could do. To varying degrees, these government policies tackle procurement issues, ethical use, training requirements, transparency, the creation of use case inventories, and privacy or security concerns.

For example, Utah’s AI policy gives guidelines to agencies and ties into the Architecture Review Board, which evaluates new AI tools. “We need to understand the data privacy and security of our data,” said CIO Fuller. “It’s important to us as a state that personally identifiable information and state confidential information not go into large language learning models that would be making that data public in any way.”

The policy also makes it clear that employees who use generative AI to help compose a document are still fully responsible for the document. “You can’t blame an AI tool if you say something silly with it. So, we’re trying, by policy, to try to enforce some of these things,” Fuller said.

A proliferation of executive orders

In the late summer and fall of 2023, states began to jump forward with the formation of advisory groups or task forces, such as one that was focused on artificial intelligence and the workforce, which was signed by [Wisconsin Governor Tony Evers](#) on August 23, 2023. The following use case describes how Wisconsin used GenAI to speed up its tax process.

7. The first public announcement of the power of generative artificial intelligence came with OpenAI’s ChatGPT release in November 2022.

WISCONSIN USE CASE

Speeding the tax process

The State of Wisconsin's Department of Revenue (DOR) won a [Federation of Tax Administrators \(FTA\) award in 2025](#) for a project intended to "modernize outdated computer systems and process tax documents and inbound mail systems," according to an FTA release.

Wisconsin used generative AI-based image recognition technology to securely get to a higher state of accuracy on scanned tax forms. This includes not just the income tax, but any kind of document scanned by the state's revenue department.

Wisconsin spent February through June converting all inbound mail. By the end of June, the state was able to switch over to process three to four hundred thousand paper items. The department plans to expand the system within the next two years to handle over seven million pages of data being captured annually.

Before using AI in this way, Wisconsin followed standard practice in capturing data from papers. High resolution images were run through software, which did traditional optical character recognition. That process had a relatively high error rate for typewritten documents (different fonts, for example, could easily confuse the software). Additionally, by the time the documents went through various stages their clarity was reduced. "It's kind of a photocopy of a photocopy problem," said Ryan Minnick, FTA chief operating officer.

And it was somewhat less accurate for handwritten documents, in which, for example, it was easy to confuse the number 7 with the number 1.

The new AI-based technology is paying off in several ways. It showed a 300 percent per hour increase in document processing and a 56 percent savings in labor costs, which could allow overworked staff members to focus on other necessary tasks. according to Keith Gross, section chief of *Wisconsin* DOR's Division of Technology Services.

The model has been trained to be very accurate. Initially, people are still verifying it, but as the model progresses, it improves. Any kind of generative AI-based image recognition tool is not only returning the data, but it's also indicating how confident it is based on all the parameters that are set. If the tool says, 'I'm 98 percent sure I got this right,' and the user agrees that it's 98 percent accurate or 99 percent accurate, the user can start to rely on the automation a little bit more. Instead of sending every single image to a person to validate, fewer go for manual review.

Although the system has had a great deal of success in producing accurate outcomes, the Department of Revenue recognizes that there needs to be ample human intervention. In fact, said Wisconsin's Gross, "in the early days of using the system, we frequently found cases where the AI-captured value was correct and the verifier altered it to an incorrect value. So, even though everybody says human is the gold standard, we always knew that humans made mistakes too."

With that in mind, Gross said, "What we'll probably be doing in the near future is another study to ask, 'Are we reaching the point where a correction from a verifier is more often making the data more or less correct?' And if they make it incorrect more often that might be the point where we realize we need to cut back human review and make them only look at the stuff that is clearly questionable based on rules. I don't know that we're there yet. The early ones we felt were heavily influenced by people learning and adapting to the system."

At that point, the basics of GenAI governance started to appear in multiple gubernatorial executive orders. For example, on September 6, 2023, [an executive order by California's Governor Gavin Newsom](#) noted key business challenges that the state faced and opened up opportunities to vendors to be innovative "based on where artificial intelligence generative AI was and how it could help solve state problems."

Jonathan Porat, California's chief technology officer, pointed out that the 2023 executive order is specifically related to GenAI. The use of artificial intelligence in governments at every level has existed for many years, but generative AI opened new ways to think about the potential uses and enabled a different level of understanding for a wider group of people. "Generative AI is giving us the ability to take on challenges that traditional technologies could never fully address," said California's Bailey-Crimmins.

Two other executive orders followed, both on September 20, 2023 from [Virginia Governor Glenn Youngkin](#) and [Pennsylvania Governor Josh Shapiro](#). As the months passed, multiple executive actions were taken in a variety of other states, including Maryland, Massachusetts, New Jersey, Ohio, Oregon, Rhode Island, and Washington.

Most recently, [North Carolina's Governor Josh Stein](#) signed an executive order on September 2, 2025, "Advancing Trustworthy Artificial Intelligence."⁸

States have continued to focus legislation and executive actions on the topic through 2025, as task force reports have emerged, and new concerns and opportunities bloomed. Among the topics receiving attention are AI use inventories, the handling of higher-risk endeavors, privacy and security protections, and the practices affecting the interplay between central governments and agencies.

8. Prior to Governor Stein's September 2025 executive order, the North Carolina Department of Information Technology had established principles, practices and guidance to state agencies for a "Government Responsible Use of Artificial Intelligence Framework" in August 2024.

As described earlier in this report, on July 11, 2025, [Virginia Governor Youngkin signed an executive order](#) focused on the use of agentic AI for reducing regulations. Citing the challenges of “job displacement, social and privacy intrusion, and the potential for both positive and negative impacts on education,” the Virginia executive order directed the Office of Regulatory Management, the chief information officer and “relevant secretariats to review existing laws and regulations.” It also ensured that state government’s use of AI be “transparent, secure and impartial” with safeguards to “protect individual privacy. An emphasis was placed on ethical and legal standards, as well as “policy standards for when and how AI can be used across state government.”

In addition to executive orders, governance actions also flow from central information technology or administrative offices, such as the policy that emanated from Ohio’s Department of Administrative Services in December 2023, labeled [Use of Artificial Intelligence in State of Ohio Solutions](#).

Legislators in many states have also introduced multiple artificial intelligence bills in recent years, although many have not passed or are still pending, according to the National Council of State Legislatures 2025 [artificial intelligence database](#).⁹

While some state legislatures, for example those in Maryland and Texas, have focused on internal governance practices, the focus of state legislation has more generally aimed at “protection for citizens” than at internal state governance questions. According to the authors of an August 18 article from the Brookings Institution, “At present, the states see the misuse of AI as something that citizens need to be protected against versus the appropriate use of AI as an opportunity for better services to citizens.”¹⁰

A spectrum of governance approaches

Only time will tell what works best for AI governance, but right now a wide variety of approaches exists. Two states that fall at extremes in terms of the control they have over individual agency decisions are Ohio and Texas, both of which are making advances in their use of AI.

Ohio has an AI Council structure, which is outlined in its AI policy, as is the governance process in which agencies submit generative AI use case ideas to the Council for state approval and monitoring.

In that state, any AI solution must be fully vetted within an agency before it even comes to the AI council, according to Vanitha Zacharias, who leads the state’s AI and AI governance efforts and heads the Investment and Governance Division within the Office of Information and Technology. After the governance process is solidified within an agency, it submits its use case to the AI council and fills out a spreadsheet that requests needed information. “That solidifies the governance process within the agencies across the state,” Zacharias said. “We are building the governance layers from the bottom up within the agencies to the state leadership.”

9. Artificial Intelligence Legislation Database, National Conference of State Legislatures, <https://www.ncsl.org/financial-services/artificial-intelligence-legislation-database#>.

10. Gregory S. Dawson, Kevin C. Desouza, James S. Denford and Marc Barda Picavet, “How different states are approaching AI,” Commentary, Brookings, August 8, 2025, <https://www.brookings.edu/articles/how-different-states-are-approaching-ai/>. Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence.



After approval by the AI Council, the agency team is provided with a list of approval requirements—actions that need to be followed as the idea moves forward. As of September, by the time an agency's idea made it through the process out of a total 125 approved use cases, 1 use case was canceled, 3 use cases were withdrawn, and 5 use cases did not proceed after proof of concept was successfully completed. About half are currently in production.

By contrast, Texas CIO Amanda Crawford said that while agencies need approval for some tech issues in Texas, that's not the case for AI. There are central state requirements for the certifications that are needed for cloud usage and a requirement to purchase off of the state's cooperative contracts, unless an exemption is given by the state.

As CIO Crawford said, in Texas there are “guidelines, guardrails, and frameworks that everyone will comply with, but as far as the central IT agency saying thumbs up, thumbs down to AI use cases or deployments by state agencies—No.”

While it's possible this may create more vulnerability in an extremely decentralized state, it also could move AI forward more quickly. And while Texas has not implemented central control over use cases, state leadership is ensuring that AI deployments include human oversight and review, particularly when the AI results in consequential decision-making. Texas is also currently working on an AI code of ethics draft, adopting NIST minimum standards for AI risk management, standardized disclosure forms and notices about public-facing AI.

Ongoing governance issues

A variety of additional issues concerning the newest uses of artificial intelligence are sparking concerns and interest. A few of the issues that governments are grappling with follow.

- **Problems with enforcement.** Alan Shark described many of the governance policies he's observed at both the state and local level as looking like the Ten Commandments, with a "thou shalt not do this" tone. "One of the biggest problems with all those that I've reviewed is they often leave out any kind of penalty," he said. The missing information is that if you do what's forbidden, what happens?
- **Shortcomings in data governance.** Data governance has a strong relationship with AI governance and in many cases, it's "data governance that needs work," said Shark. "AI lives and feeds on data and if data is not properly classified and, in some cases, redacted, that could cause a lot of problems."

Data governance also was a pivotal area of concern to NASCIO's Robinson. "Data governance is a critical factor in the use of AI. If you have these closed, large language models and you're ingesting terabytes of data from state agency X, or from multiple state agencies, you've got to be concerned about the data quality."

- **Measurement and evaluation of AI.** Utah plans to do surveys that track how much time people use the tool and how much time they believe it saves them. In a pilot survey, heavy users told the state that their use of AI was saving them more than three hours a week, though Fuller emphasized that this was a preliminary look. His major point is that while the state feels it is "pretty good" at measuring adoption, measuring impact is hard and still a work in progress. "I think everybody's struggling with measuring the actual impact of GenAI tools," he said.
- **Integration with other technology issues, external GenAI products, and state management practices.** NASCIO's Robinson discussed how general trends affecting the strategic role of the CIO office also affect how AI is handled. He said many AI elements mirror and flow out of what's in place in states for the performance management of their technology and their more general governance model. AI does not happen in a vacuum and its governance is also affected by whether states have well-defined risk management, are focused on business transformation or business process improvement, and whether they have a shared decision-making process with agencies and Enterprise Portfolio Management practices.

Mark Raymond, CIO in Connecticut, noted that as generative AI took off, vendors began adding AI elements to existing products. "The industry was rushing to add AI into the environment," he said. "Some of it was outside of our control and we needed to account for it."



Observations from Experts

The research for this report included careful review of a number of documents and reports. But its major findings were uncovered through a series of interviews with CIOs and high-level technology officials in over a dozen states spread across the nation—from Connecticut to California, with many stops in between.

Over the course of these interviews, some clear observations emerged:

- While there's a level of fear about humans being replaced by AI, the most commonly repeated statement in interviews reflected the vital importance of keeping "humans in the loop."
- Though GenAI has been getting significant attention, many states are still in the early stages of implementation. Pilots abound, but taking the programs to scale is less common than might have been anticipated.
- It's still unclear as to what impact AI will have on the number of employees on state payrolls—though strong arguments can be made that at least some specific jobs will change.
- Though most states are aware that they need governance over their agencies' use of AI, the approaches vary from very constrictive controls to greater flexibility.
- The most common use of GenAI currently is in a variety of chatbots, which can help to answer questions at lightning speed. However, so far much of the sophisticated use of these tools is for internal purposes and not outward facing for public use.

Conclusion: Vital Next Steps

With technology in general and government practices in a transitional period, this report covers the state of AI in the states with a point in time analysis for the fall of 2025. A number of principles will contribute to the successful use of AI while avoiding pitfalls. They include:

- Be certain that existing processes used in government—like gathering accurate data—function well, before utilizing them to feed AI systems.
- Put the time and money necessary into adequate training for government employees at all levels—and the public as well.
- Whatever governance structures are put into place, make certain that they are being followed, with the necessary funding available to do so.
- Ensure genuine transparency before embarking on an AI project to preserve accountability.
- Evaluate the benefits of AI. This goes beyond simply measuring speed and efficiency, and extends into measurements of responsible use, and the trust of the populace that AI is serving them well.
- Be clear about what AI cannot do (at least yet). Extending its use in ways that don't add value can present a risk at a time some view AI as a panacea to a whole host of government problems.

This is by necessity just a sampling of State efforts while they advance the use of AI. As time passes AI will become more widely used, including in ways not imagined now. With that in mind, it's vital that states continue to observe the unpredicted consequences and the unpredicted benefits of AI, and carefully consider the lessons learned from them.



About the Authors



Katherine Barrett and Richard Greene

Principals

Barrett and Greene, Inc.

P: 212-684-5687

E: greenebarrett@gmail.com

Katherine Barrett and **Richard Greene**, who have analyzed researched and written about state and local government for over thirty years are visiting fellows at the IBM Center for The Business of Government; contributors to *Governing Magazine*, executive advisors for the American Society for Public Administration; editors and administrators of their own website “dedicated to state and local government” at greenebarrett.com; advisors to the Government Finance Officers Association; columnists for *Government Finance Review*; commentary editors for the *International Journal of Public Administrators*; and senior advisors to the Government Finance Research Center at the University of Illinois, Chicago. They are also fellows at the National Academy of Public Administration. Their latest book, *The Little Guide to Writing for Impact*, was published in 2024.

Recent Reports from the IBM Center for The Business of Government



GenAI and the Future of Government Work

by William G. Resh



Embedding Strategic Foresight into Strategic Planning and Management

by Bert George



Resilience in action: Crisis leadership through innovation, collaboration, and human-centered solutions

by Julia Carboni




Leadership Framework for an Agile Government

by Pallavi Awasthi and Kuang-Ting Tai



The Opportunity Project

by Joel Gurin and Matt Rumsey



Drought, deluge, and data: Success stories on emergency preparedness and response

by Samanta Varela Castro, PhD, Keri K. Stephens, PhD, Jongeun You, PhD, and Wendy D. Chen, PhD



IBM Center Research Announcement

by Dan Chenok



How TECHNOLOGY Can Drive GOVERNMENT EFFICIENCY

by Dan Chenok



For a full listing of our reports, visit businessofgovernment.org/reports

About the IBM Center for The Business of Government

Through research stipends and events, the IBM Center for The Business of Government stimulates research and facilitates discussion of new approaches to improving the effectiveness of government at the federal, state, local, and international levels.

About IBM Consulting

With consultants and professional staff in more than 160 countries globally, IBM Consulting is the world's largest consulting services organization. IBM Consulting provides clients with business process and industry expertise, a deep understanding of technology solutions that address specific industry issues, and the ability to design, build, and run those solutions in a way that delivers bottom-line value. To learn more visit ibm.com.

For more information:

Daniel J. Chenok

Executive Director

IBM Center for The Business of Government

600 14th Street NW

Second Floor

Washington, D.C. 20005

(202) 551-9342

website: www.businessofgovernment.org

e-mail: businessofgovernment@us.ibm.com

