GOVERN FOR AMERICA

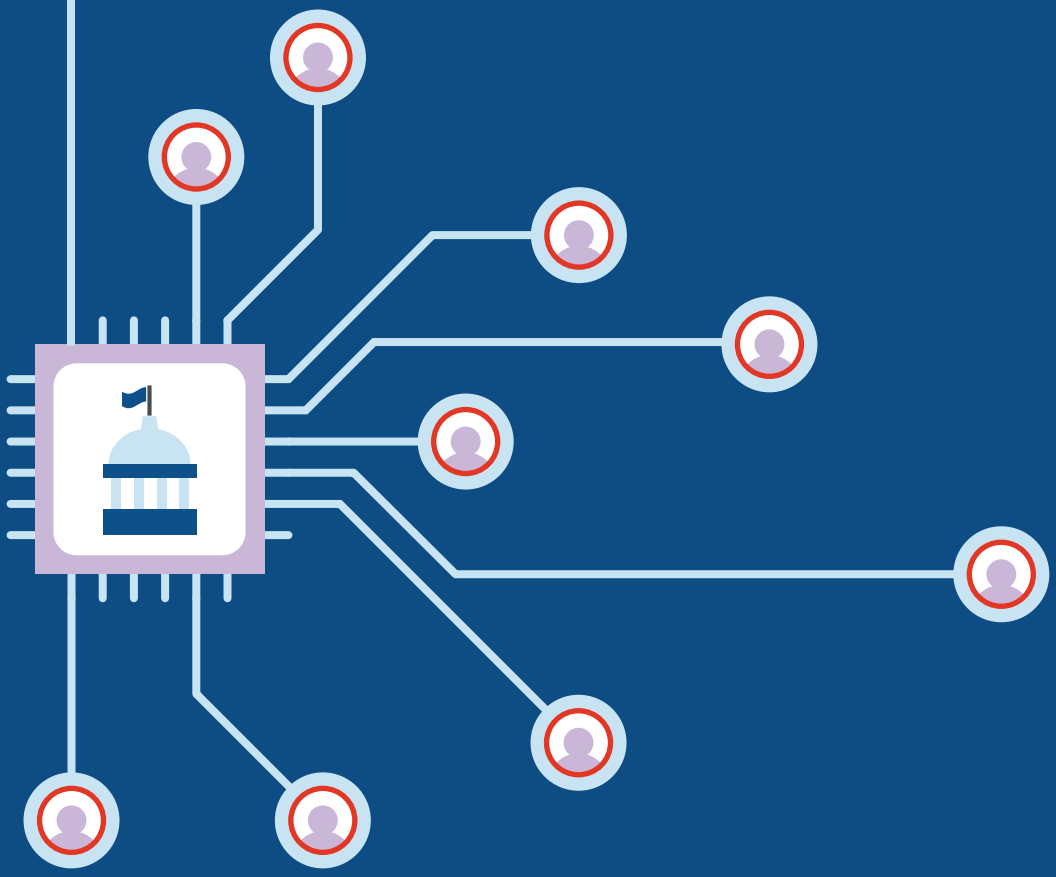

AI x Talent

A Roadmap for Building, Staffing and Scaling AI Capacity in State Government

April 2025

The Team

The AI x Talent project was led by Govern For America. We are incredibly grateful to the 50+ government, non-profit, and university stakeholders we interviewed throughout our research. We would also like to thank The Public Interest Technology University Network and The Center for Public Sector AI, for acting as convening partners.

About GFA

Govern For America (GFA) is a non-profit founded in 2018, with the mission of building the next generation of public servants to create a more responsive government that better reflects and serves our communities. GFA builds a pipeline into public service by connecting high-potential recent college and post-grads to high-impact roles in state and local governments through a two-year fellowship program. GFA recruits, trains, and places the graduates, activating a network of civic leaders across the nation. Since its founding, GFA has placed nearly 100 Fellows who have impacted 47 million people and driven $3.7B in funding to communities through their work. GFA Fellows have been placed in agencies including: Department of Health Services (Arizona), Office of Career and Technical Education (Michigan), Department of Service and Civic Innovation (Maryland), and Department of Early Childhood (Colorado).

Led by Govern for America
Octavia Abell, Founder and CEO
Talita Elizeu, Director of Partnerships

With research support from
Cara Eckholm, Research Director
Brandon Luebbehusen, Research Manager
Ben Oldenburg, Graphic Design
Andrew Eckholm, Editor

This research was made possible thanks to philanthropic grants from Microsoft and the Public Interest Tech Fund, a pooled vehicle of the Ford Foundation, Siegel Family Endowment, Patrick J. McGovern Foundation, and Pivotal Ventures. GFA retained full editorial independence over the research process, findings, and report publication.

Contents

Executive Summary

Since the launch of ChatGPT in November, 2022, at least thirty-two U.S. states have formed Artificial Intelligence (AI) task forces and Governors routinely list AI among their top priorities. Yet there has been limited research published on how states are—or should be—using AI to improve their own internal staff capacity, and delivery of constituent services.

From September 2024 to March 2025, we conducted semi-structured interviews with a cohort of state government IT leaders and other tech ecosystem experts to help fill this gap. **Our goal was to document how leading states have rolled out AI applications across core government functions, capturing what worked—and what did not—in an actionable roadmap for states that are earlier in their AI journey**. Through interviews with university career advisors, we also explored how states can attract and empower AI-native talent from the next generation, to help spread AI efforts across government in the years ahead.

Our findings suggested six steps (Figure 1) that state CIOs generally recommend following when introducing AI into government:

1. **Determine Ownership**
2. **Issue Guidelines**
3. **Raise Awareness**
4. **Source Use-Cases**
5. **Launch Off-the-Shelf Solutions**
6. **Prototype Custom Solutions**

These six steps are not necessarily sequential, and may be–often should be–repeated, as the technology advances. For example, interviewees emphasized the need to keep responsible use guidelines as living documents, with some states issuing updates on a bi-annual basis. While few states have hired new employees for AI-specific roles yet, nearly all agree on the importance of empowering someone to act as the point person on AI. As of this report's publication, the most advanced states have just finished moving through these initial six steps, and are now submitting budget requests to staff and scale their AI efforts.

> "...nearly all [states] agree on the importance of empowering someone to act as the point person on AI."

Figure 1.

Six Steps to Developing an AI Strategy

The following six steps, recommendations, and resources were contributed by the cohort of thirteen leading states that participated in research for our study. They reflect learnings from the first two years of these state's AI journeys.

	Step	Recommendations	Best Practices & Resources
	**Determine ownership of AI strategy**	• Establish an AI "Center of Excellence" within the IT Department to work with agencies • Staff with an AI director, product manager, and policy manager	Maryland created a CAIO position staffed with two supporting roles in policy and product management
	**Issue guidelines for responsible use**	• Issue guidelines covering acceptable uses, data protection and privacy, AI bias prevention, and procurement • Create a process for iteratively update the guidelines on at least an annual basis	Vermont has both a Code of Ethics and State Guidelines
	**Raise awareness across agencies**	• At a minimum, roll out responsible use guidelines for awareness, and empower managers to make informed judgement calls • Ideally, roll out an optional AI training program that helps staff understand the limits and potential of AI – and gets them excited!	New Jersey worked with Innovate U.S. to develop free training modules
	**Source use cases and prioritize them**	• "Push" standard use cases to agencies (e.g., tools for content creation, coding assistance, language translation, internal "policy bots") • "Pull" ideas from agencies based on their priorities, and then prioritize them by confirming they are both responsible and productive uses	Connecticut has an "AI Intake Form" for agencies to submit use cases for review by a central team
	**Launch standard off-the-shelf solutions**	• Start by deploying "off-the-shelf" tools (e.g., OpenAI's ChatGPT, Microsoft's Copilot) for internal-facing use cases • Create a sandbox for a subset of employees to test tools before scaling to the broader organization	Pennsylvania launched a ChatGPT pilot with 175 employees and collected structured feedback
	**Prototype custom point solutions**	• If building custom point solutions, consider starting with simple prototypes – AI pilots can experience up to a 50% failure rate • If buying, run a structured procurement process that enables a pilot with multiple vendors, at low cost	California issued Requests for Innovative Ideas for five priority use cases

> "...one of the most important factors in a given AI pilot's success was structured, well-labeled data."

Leading states have reported impressive results from pilot projects. Commercially available, generative AI tools for writing, coding, and editing are already saving time and improving quality. In Pennsylvania, for example, ChatGPT enterprise saved users more than 90 minutes per day, while Vermont has reduced technical debt through automated code review. While most state IT leaders are still reluctant to have constituents interact directly with AI agents, internal-facing chatbots to assist call center and other agency staff in answering constituent questions are becoming common. At the same time, State Chief AI Officers (CAIOs) also caution that many AI pilots will fail at first—and that learning from failure, and iteratively training AI models, is part of the experimentation process. Generally, one of the most important factors in a given AI pilot's success was structured, well-labeled data. Our research strongly suggests that states lacking in data science capabilities will want to address this deficit before embarking on an AI journey in earnest.

Most states envisioned a centralized, Center of Excellence (COE) approach to AI, with the COE headed by an AI lead and assisted by a policy manager (covering responsible use) and a product manager (covering productive use). While these roles require talent experienced in large organizations and IT functions, many interviewees also saw an opportunity for early-career talent to fill AI and AI-enabling roles. Contrary to popular perception, there is an excess supply of young, technically-minded talent interested in working in government. The problem is that matchmaking is not occurring effectively—a problem that structured partnerships between state governments and universities could quickly help address. Given the looming retirement of baby boomer employees, AI presents an opportunity to attract the next generation of government IT staff through impactful projects that move much quicker than typical government initiatives.

Philanthropy also has a role to play in accelerating state AI journeys. Leading states have just now submitted budget requests to expand their AI teams—but many CIOs anticipated multiple budget cycles before these roles are granted. A structured, philanthropically funded fellowship program could help bridge this funding gap, allowing leading states to continue to build on their wins, and provide tools and talent for states just beginning to grapple with AI. A fellowship could also be paired with a community of practice: given that all states perform the same basic set of functions and are therefore experimenting with a similar set of AI use cases, there is ample opportunity for cross-state learning and A/B testing.

> "Contrary to popular perception, there is an *excess* supply of young, technically-minded talent interested in working in government. The problem is that matchmaking is not occurring effectively."

Background

The Artificial Intelligence Opportunity in Government

In November 2022, OpenAI launched ChatGPT, inaugurating the era of generative artificial intelligence (AI). Experts have said that AI will automate 2.4 million jobs by 2030—but they have also projected that AI will create trillions of dollars in economic value, in part through enhanced productivity. A core question underlying the growth of the technology is how to reskill workers so that they can participate in the emerging AI economy.

In the United States, 14% of the labor force works for the federal, state or local government. Government officials across the country are navigating what the emergence of AI might mean for their own workforce, and how they deliver services to America's more than 330 million residents. Much like their counterparts in the private sector, government Chief Information Officers (CIOs) must figure out how to enable their workforce to use AI responsibly and productively.

Under the Biden Administration, the Federal Government launched an "AI Talent Surge," aiming to recruit talent for hundreds of new "AI roles" across federal agencies in 2024 and 2025. The governors of at least thirty two states launched AI task forces (or similar initiatives) to help position their states as leaders in the field. Some of the country's better-resourced cities launched their own AI experiments, though the results have been mixed, demonstrating both the opportunity and current limits of the technology. For example, San Jose has had success using AI to monitor traffic safety and improve bus travel times by implementing dynamic signal prioritization. By contrast, an AI chatbot created by New York City to help small business owners navigate regulations doled out incorrect—and at times even illegal—advice.

Case Study: The Federal Government's AI Talent Surge

In October 2023, former President Biden signed an executive order to promote the responsible use of AI, directing Federal agencies to act on priorities related to consumer privacy and protection, AI innovation and competition, and America's global AI leadership. Given the AI expertise required to act on these priorities, the executive order also created the National AI Talent Surge to support agencies in the planned hiring of more than three thousand AI and AI-enabling individuals from April 2024 to September 2025. The majority of this planned hiring was for the Department of Defense. The Trump Administration has at least notionally indicated that hiring and preserving AI talent continues to be a priority.

The reality is that many parts of government are still hiring foundational data teams. Deploying Large Language Models (LLMs) at scale will likely require substantial advances in state capacity. However, the surge in interest in AI has created momentum to invest in underlying data science capabilities, a critical first step in any government's AI journey. According to one agency head, **"Even if AI is a bit of a trojan horse, it does free up space to have a conversation about the government's needs for a technology workforce more broadly."**

Definitions

Data Science is the process of collecting data from multiple sources, organizing the information, translating the information into insights, and communicating those insights.

Artificial Intelligence (AI) is the ability of a machine to simulate human intelligence by making logical conclusions based on data and information.

Generative Artificial Intelligence is the subset of AI tools that use existing data sets to create original content like documents, images, and code, or to suggest edits. Generative AI was the primary focus of our study.

Effective AI solutions rely on robust data sets—but for many government agencies, data sources are siloed today. Data scientists are needed to aggregate and clean data so that it can serve as an input for AI applications. One agency head summarized this critical link by saying, **"We went through the mapping [of necessary AI-related roles], and there were a number of data scientists. That is something that's not fully built out at every department...That is definitely going to be a huge contingency on the success of us rolling out AI."**

Even though deploying LLMs at scale to improve constituent services still feels far away for some, there is an eagerness to get started today, especially with commercial, off-the-shelf generative AI solutions. As one official put it, **"When we rolled out cyber training, I had to force it on people, but people are beating on my door for access to tools like AI transcription and ChatGPT."**

More broadly, given the rapid progress in AI technology, government IT officials are optimistic that AI will soon help address a wide range of public sector problems, ranging from automating fraud detection to more efficiently dispersing natural disaster assistance. AI's value is already on display in first-mover states like Vermont, which began developing an AI strategy in 2018, and now develops and deploys between ten and fifteen new AI tools annually.

How States Are Responding

Much has been written about the impact and opportunity for AI in the federal government. However, there has been limited research about the opportunities and limitations of AI in State and local government—despite most Americans interacting more with their state and local government than federal government for services. We sought to fill this gap with the AI x Talent Project, informed by qualitative research conducted with State Chief AI Officers (CAIOs) and equivalent officials across the country.[1]

From California to Massachusetts, governors have listed AI as one of their state's top priorities. When all 50 State CIOs (or their equivalents) were surveyed, they ranked AI as their second top priority, only behind cybersecurity. Partly in response to this momentum, at least 32 U.S. states have launched task forces (sometimes referred to as "working groups," "advisory councils," or "steering committees") to help refine their state's AI strategy (Figure 2), typically consisting of representation from the Governor's Office, the state's IT office, and AI researchers from local universities. While most task forces were created by executive order and run by

1 While cities also play a significant role in service delivery, few have sufficiently large technology teams and resources to have dedicated AI strategies. Some exceptions include New York, Chicago, San Francisco, San Jose, and Boston.

AI x Talent Research Scope

In-Scope

Domains typically handled by the IT Department

1 Responsible Use

- How do we ensure that AI reflects our values and avoids algorithmic bias?
- How do we implement AI while keeping the data it relies on private and safe?

2 Productive Use

- How do we use AI to enhance service delivery to our constituents?
- How do we reskill agency staff to more effectively use AI in their jobs?

Out-of-Scope

Domains typically handled by a state's labor or economic development department

3 Business Attraction

- How do we attract, grow and retain AI companies in our state?

4 Workforce Development

- How do we upskill workers in AI, while minimizing risk of job displacement?

the governor's office, legislatures in at least eight states created and convened their own task forces.

The reality is that an "AI strategy" can mean many different things, and depending on the state, tasks forces tend to cover some combination of the following four distinct yet complementary aims: Responsible Use, Productive Use, Business Attraction, and Workforce Development.

While all four topic areas are of critical importance to a state's economy, the AI x Talent project study specifically focuses on the responsible and productive use of AI by state governments—domains typically covered by a state's Information Technology agency. We then looked at the talent needs required to deliver on responsible and productive use effectively.

Can AI Help Solve Chronic Worker Shortages?

While studying potential adverse workforce impacts (i.e., job displacement) was not an explicit goal of our study, we heard indirectly about this subject in our interviews, particularly in rural states with more limited government capacity. Contrary to the narrative that AI is likely to result in job displacement, there was hope in more rural parts of the country that AI may help to alleviate chronic staffing shortages, which are likely to worsen in the coming years. Over 42 percent of the federal workforce is over the age of 50, compared to just 33 percent for the U.S. workforce overall. Many state governments face similar ratios.

FIgure 2.

States with AI Task Forces

(as of February 1, 2025)

Responsible and productive use of AI—the subject of our study—is a priority for nearly all of the 32 task forces launched by state governments. While many states have created AI task forces, often staffed by volunteers, fewer have yet moved to empower an AI lead or create a permanent AI Center of Excellence (CoE) in the state's IT function

32 states have established AI task forces

At least **13 states** have an AI lead who participated in our research

State	Task Force	Launched	In-Scope for Research: Responsible Use	In-Scope for Research: Productive Use	Out-of-Scope for Research: Workforce and/or Economic Development
Alabama	Y	2024	●	●	
Alaska	N				
Arizona	Y	2024	●	●	
Arkansas	Y	2024	●	●	
California	Y	2023	●	●	
Colorado	Y	2024	●		
Connecticut	Y	2023	●		●
Delaware	Y	2024	●	●	
Florida	N				
Georgia	Y	2024	●	●	
Hawaii	N				
Idaho	Y	2024	●	●	
Illinois	Y	2023	●		●
Indiana	Y	2024	●		
Iowa	N				
Kansas	N				
Kentucky	Y	2024	●	●	
Louisiana	N				
Maine	Y	2024	●	●	●
Maryland	Y	2024	●	●	
Massachusetts	Y	2024	●	●	●
Michigan	N				
Minnesota	Y	2023	●	●	
Mississippi	N				
Missouri	N				
Montana	N				
Nebraska	N				
Nevada	Y	2024	●	●	
New Hampshire	N				
New Jersey	Y	2023	●	●	●
New Mexico	N				
New York	Y	2024	●	●	●
North Carolina	N				
North Dakota	N				
Ohio	Y	2024	●	●	
Oklahoma	Y	2023	●	●	
Oregon	Y	2023	●	●	
Pennsylvania	Y	2023	●	●	
Rhode Island	Y	2024	●	●	●
South Carolina	Y	2024	●	●	
South Dakota	N				
Tennessee	Y	2024	●	●	
Texas	Y	2023	●	●	
Utah	N				
Vermont	Y	2018	●	●	●
Virginia	Y	2024	●	●	
Washington	Y	2024	●		
West Virginia	Y	2024	●	●	●
Wisconsin	Y	2023			●
Wyoming	N				

Research Framing

The AI x Talent Project

One of the first actions taken by many states after the November 2022 launch of ChatGPT was to enact a responsible use policy governing the use of generative AI. As early-adopter states now move from policy into piloting and deploying AI models, they are learning about the most productive uses, and the limitations, of the new technology. Critically, this transition from policy to deployment also correlates with a need to scale up internal staffing capabilities.

IT responsibilities and institutions differ across the fifty states. For example, some state governments have a centralized IT department, while others have IT sitting as a function within operating agencies. But at their core, state governments are responsible for delivering the same basic set of services to their constituents. They all provide social services like unemployment benefits, build and maintain roadways, and manage healthcare systems—tasks that may or may not be enhanced by AI. These common activities make states a particularly fruitful sample set for comparing AI journeys, and lead to ample opportunities for cross-state learning.

To that end, GFA decided to conduct a foundational research study with three main objectives, detailed to the right.

Research Objectives

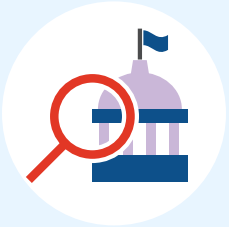

Assess the degree of preparedness within states to deploy AI, and detail the most successful deployments illustrated by real examples from leading states

Profile demand from early-career talent to fill state technology roles to enable AI deployment, and identify the challenges and opportunities associated with placing talent

Develop a working model for how state governments can utilize fellowships or other innovative staffing mechanisms to begin expanding their AI capabilities

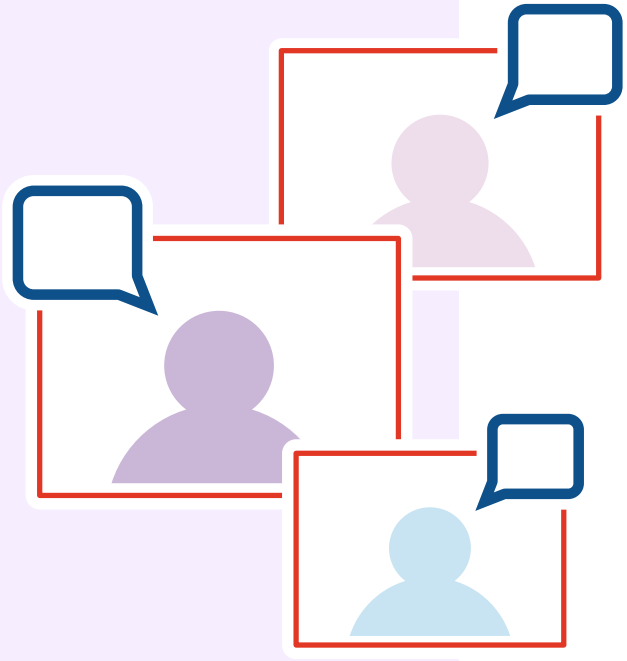

Methodology

We completed our study over a six-month period from September 2024 to March 2025, conducting research and interviews with three distinct but complementary groups: State Governments, Tech Talent Experts, and University Career Advisors. The quotes included throughout the report are derived from these interviews. All quotes were anonymized to ensure respondents felt comfortable speaking freely.

Respondent Profiles

1. **State Governments**: 45-minute, 1-1 interviews, and a focus group with 13 offices of state Chief Artificial Intelligence Officers, state Chief Information Officers, or equivalent offices responsible for setting state AI agendas.[1][2]
2. **Tech Talent Experts**: 45-minute, 1-1 interviews with 15 individuals who study or run programs to develop technology capacity in government, like the Tech Talent Project and U.S. Digital Response.
3. **University Career Advisors**: Focus group with career advisors from 11 member universities of the Public Interest Technology University Network.

1 In addition to State CAIOs, we interviewed economic development organizations in four states that did not yet have a structure around who was managing the development of a local AI strategy.

2 The sample of State government responses is limited to states with a dedicated AI lead (or equivalent) in place. States lacking a formalized AI lead presented challenges for gathering qualitative input, but our analysis was augmented by publicly available information where available.

Figure 3.

Research Participants

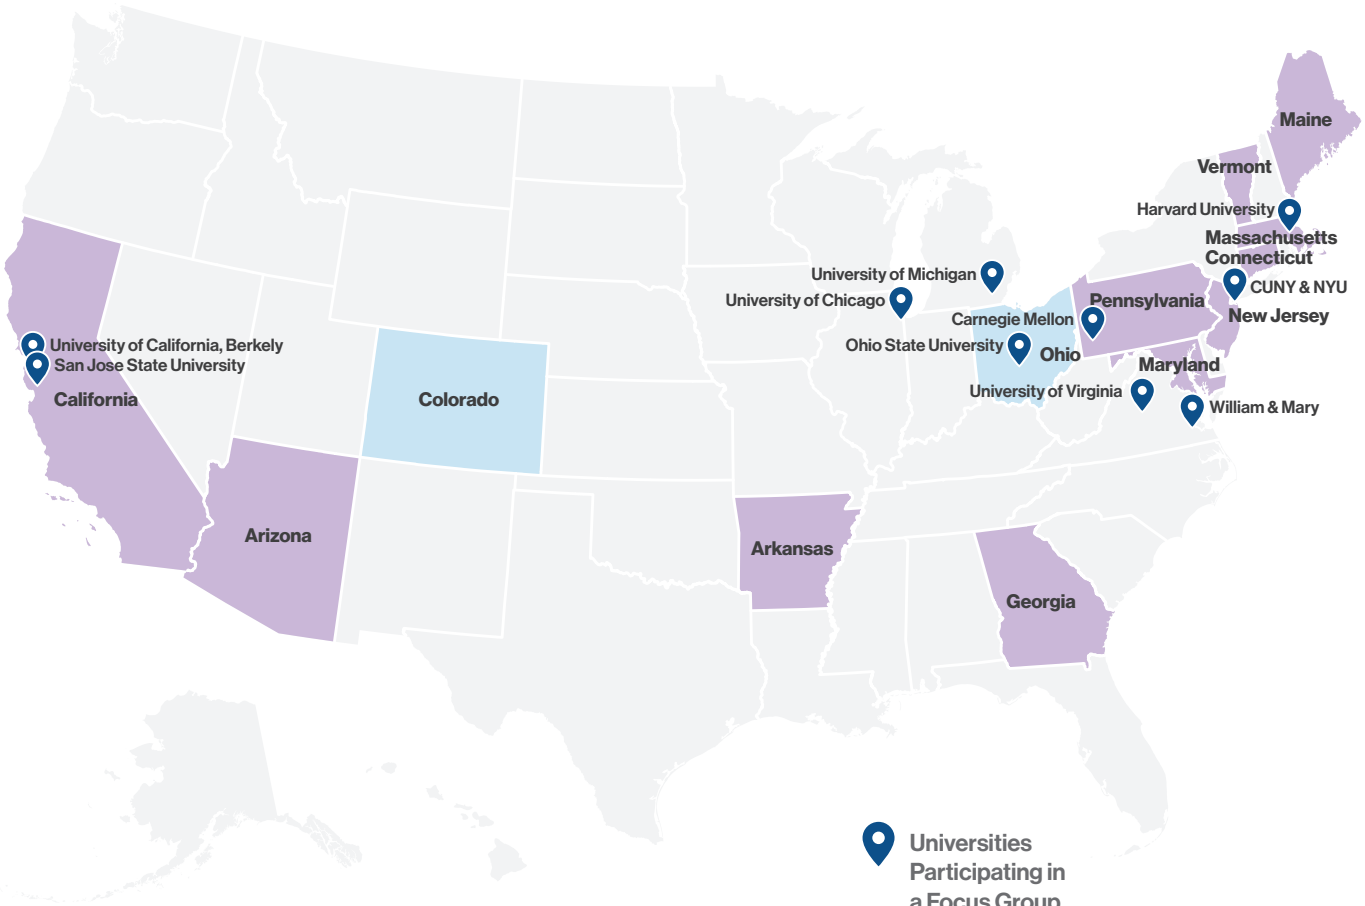

Technology Talent Experts

- Center for Public Sector AI
- City of San Jose
- Department of Homeland Security
- Ford Foundation
- Generation West Virginia
- NASCIO
- New America/PIT-UN
- National Governors Association
- Office of Personnel Management
- Renaissance Philanthropy
- Tech Talent Project
- U.S. Digital Response
- Volcker Alliance

Findings

The Six Steps to Developing an AI strategy

While each state follows its own journey in developing a framework for the responsible and productive use of AI, through our research, we attempted to extrapolate commonalities. We mapped out a six step journey (Figure 4), which tracks roughly to the first two years of a state's rollout of generative AI tools. As AI technology matures, and as states just now building out AI capacity draw on learnings from first movers, these six steps will likely take less time.

The steps are not necessarily sequential, and may be repeated—for example, some states issued an initial set of responsible use guidelines in early 2023, and have since issued successive updates. One respondent told us, "About half the things we tried didn't end up working. There is a big iteration loop." In repeating steps, State leaders used learnings from failed pilots or flawed policies to help craft more successful efforts.

Based on progress made through these first six steps, CAIOs in leading states are now trying to advance from prototyping to staffing and scaling AI efforts. Two years in, the most far-along states have just reached this inflection point, and are submitting requests for additional headcount to their legislature.

Figure 4. Developing and Refining an AI Strategy

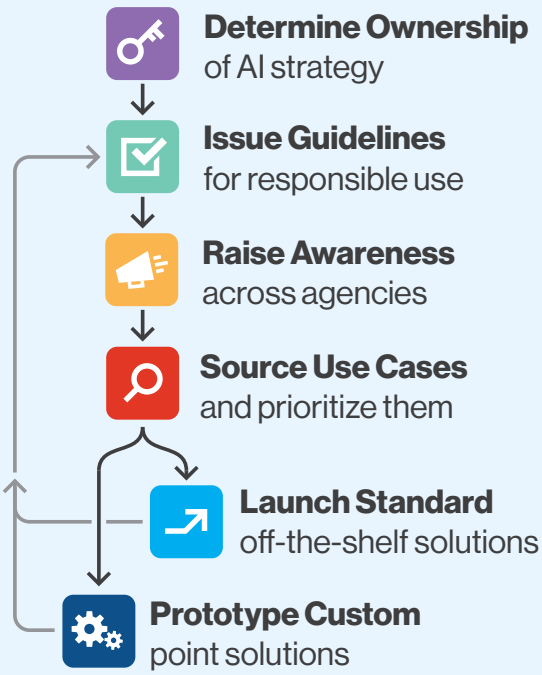

Figure 5.

Representative State AI Timelines

While conducting interviews, we asked each state to share how their AI journey mapped onto the six steps, as well as tips and warnings for each stage. In the rest of the document, we share advice from state IT leaders on how they developed an AI strategy.

Georgia

Late 2023

Determine Ownership
Hired a Chief Digital & AI Officer within the Georgia Technology Authority.

Issue Guidelines
Released initial standards to guide the responsible use of AI, which covered topics like procurement, data privacy, ethics, and training.

Raise Awareness
Hosted the 2023 Georgia Emerging Technology Summit: AI for both public- and private-sector stakeholders to understand the future trajectory of AI and ensure its responsible use.

Launch Standard
Launched the first two AI pilots, including a resident-facing chatbot and a website security tool.

Early 2024

Issue Guidelines
Convened an AI Advisory Council to offer guidance and supervision in the use of AI among state employees.

Mid 2024

Source Use Cases
Released a Request for Qualified Contractors to pre-qualify contractors for the future development of AI solutions.

Late 2024

Raise Awareness
Engaged InnovateUS to train state employees on how to use AI to enhance public service delivery while managing risks responsibly.

Early 2025

Source Use Cases
Launched an Innovation Lab for pre-qualified contractors to collaborate with state agencies in proposing AI solutions and developing proofs of concept.

California

Late 2023

Determine Ownership

Signed an executive order to authorize and enable State departments to collaborate on a responsible process for evaluation and deployment of AI (similar to an AI task force).

Early 2024

Issue Guidelines

Released guidelines for state agencies to follow while procuring AI solutions.

Source Use Cases

Released a set of Requests for Innovative Ideas to solicit potential AI use cases for known State priorities.

Mid 2024

Prototype Custom

Signed partnership agreements with five vendors to test, iterate, and evaluate AI use cases through proofs of concept.

Raise Awareness

Began offering AI training for state employees, covering four technical domains: security, data, engineering and development, and project management.

Late 2024

Source Use Cases

Hosted the GenAI Innovator Showcase to conduct market research and explore emerging AI solutions.

Raise Awareness

Launched a public-facing website to showcase AI projects happening across the state government.

Source Use Cases

Released a second set of Requests for Innovative Ideas to solicit potential AI use cases for known State priorities.

Arizona

Early 2024

Issue Guidelines

Released a statewide policy that established guidelines for the responsible use of generative AI technologies by state employees.

Mid 2024

Determine Ownership

Created a new State Data and Analytics Office within the Department of Administration, responsible for advancing data efforts and supporting generative AI usage across the state.

Raise Awareness

Engaged InnovateUS to train state employees on how to use AI to enhance public service delivery while managing risks responsibly.

Launch Standard

Created sandboxes with Amazon Web Services, Google, and Microsoft so that state employees could test off-the-shelf AI solutions in controlled environments.

Late 2024

Source Use Cases

Began soliciting AI use cases and their expected value from state employees, for review by the Department of Administration.

Issue Guidelines

Released an updated statewide AI policy based on feedback that the Department of Administration collected on AI usage in the state.

Source Use Cases

Announced the creation of an AI Steering Committee to inform future AI deployment and use cases.

Step 1: Determine Ownership of AI strategy

Establish a Center of Excellence

Of the 13 states we interviewed, nine decided to create a "Center of Excellence" (COE) to roll out AI tools (though the particular terminology varies). The COEs were most commonly housed in the state's IT department, where IT policy and production typically sits.

It is important to distinguish a COE from a task force set up to help shape a state's approach to AI.

A task force is focused on big-picture strategic goals and creating political and staff buy-in for change. A task force is likely temporary, though some states, like Vermont, have converted their task force to a permanent advisory body. Task forces are often staffed by volunteers from civic society (e.g., academia and/or prominent local companies).

A COE is a permanent group of government staffers that sit within a state's IT department, and deal with the nuts and bolts of AI implementation. That includes setting policies, standing up pilots, and eventually rolling out generative AI tools across agencies. While at least 32 states have task forces that serve an advisory role, a far smaller number have a COE or equivalent dedicated to AI implementation.

Case Study: Vermont as a First Mover

Vermont was the first state to launch an AI task force, after it was mandated to do so by its legislature in 2018. Given the then-nascent state of the technology, the task force focused on issues such as algorithmic bias, and more broadly positioning the state to be ready for AI once the technology matured. The body met regularly between 2018 and 2020, and produced a report. That report then got picked up off the shelf in 2022, after the release of ChatGPT. Following the report's recommendations, the legislature created a permanent AI council and division, and the task force now reports out to the legislature once a year. While the State's AI efforts haven't yet resulted in any new FTEs, the Chief Data Officer was given responsibility for AI (his title changed to the Chief Data and AI Officer) and some staff were given part-time responsibilities in AI. Vermont is just now hiring for a full-time role dedicated exclusively to AI.

The COE-approach mirrors a broader trend in state government IT services: centralization. The size of the federal government means that agencies and departments can employ dedicated in-house talent and expertise. States, by contrast, are unlikely to have the resources for a dedicated director of AI or equivalent within each agency. As with the transition to cyber, AI in state government is likely best spearheaded by a centralized COE that can serve the needs of every agency.

Name a Director

Most states we interviewed did not create a new AI director role—at least not yet. In general, three main approaches were taken for where to assign responsibility for AI: Emerging Technologies, Data Science, and Digital.

Some states, like Pennsylvania, created a "Director of Emerging Technology" position, who could take on responsibility for AI programs. Other states, like Georgia, gave the responsibility to the Director of Digital. Arizona empowered the Director of Data Science. Out of the states we interviewed, only one—Maryland—had hired an AI-specific director position, though a number of others were in the process of hiring for the role. States currently lacking the data infrastructure and know-how necessary to begin implementing AI will likely wish to address these deficits before making a permanent decision over where AI responsibilities rest, or before hiring on an AI team.

Most CIOs agree on the importance of someone acting as a "Director of AI" atop the COE, whether a new hire or an existing staff member. Most also saw the need for at least two further positions to assist the director, including a policy manager (to identify responsible uses) and a product manager (to identify productive uses).

Case Study: Maryland Hires a Director of AI

In October 2023, Maryland's Governor appointed Nishant Shah as the State's first Senior Advisor for Responsible AI, charged with overseeing Maryland's AI strategy.

In late 2024, the State began hiring for two supporting roles to continue implementing the strategy: a policy manager and a product manager.

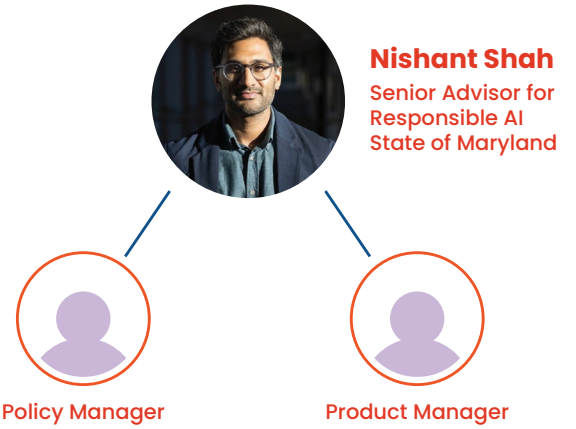

Step 2: Issue Guidelines for Responsible Use by State Agencies

Establish Guidelines

In a majority of states in our study, the executive branch (more specifically, the state IT agency) has taken the lead in creating policies to govern the use of AI. However, in at least eight states, legislatures have taken the initiative to address real or perceived shortfalls by passing bills that impact staff use of the emerging technology. As AI continues to develop and use cases in government expand, legislatures and executives may need to deepen their collaboration on creating guidelines for responsible use.

> "Your job as a head of AI is also to lead with the legislature. A more robust approach to legislation is necessary, relative to most tech roles."
>
> – State AI lead

Absent legislative intervention, however, it is generally the responsibility of the IT team to create AI policy and protocols. Typical policies cover issues such as data privacy, the potential for AI bias, and procurement policy. Most guidelines are stratified by risk. For example, an internal-facing AI chatbot would face a less stringent review process than a public-facing one.

Iterate on Policy

Because of the fast pace of development in AI, state IT leaders noted that policy needs to adapt and change for adoption to succeed. This is one reason why IT teams are often eager to take the lead in defining guidelines for AI usage rather than wait for hard-to-change legislative intervention. Many states have had to evolve their policy on an annual (or even more frequent) cadence as they build AI capabilities. Generally, the preference of CIOs has been to get an early "stake in the ground" on policy and iterate on guidelines as prototypes are built. State CIOs recommended using a standard set of AI definitions published by the Advanced Technology Academic Research Center (ATARC), for increased interoperability.

> "The guidelines are definitely living documents. We are evolving every six months."
>
> – State AI lead

Case Study: Vermont's Framework for Evaluating Generative AI Use-Cases

In 2022, Vermont's AI Council drafted a Code of Ethics governing the use of AI in the State government. However, the release of ChatGPT shortly after meant that the Council had to quickly begin iterating. Building on the initial code, Vermont adopted guidelines on staff use of generative AI for personal productivity in early 2023. These guidelines have been updated

twice since, with the latest refresh allowing government workers to use AI as part of work processes.

Vermont's guidelines on generative AI, shown below, categorizes potential use cases according to two axes: breadth of distribution (who will see the work product?) and intensity of AI use (how much of the content is being generated by AI?). As Josiah Raiche, Vermont's Chief Data and AI officer told us, "As time goes on, we expect to see more checkboxes move up into the right as people come to expect AI." That is, as Vermont's AI journey continues, the State is expecting to permit more public-facing documents to make greater use of AI.

Case Study: Arizona Iterates on Policy

Arizona released its initial AI guidelines in March 2024. However, the State's learnings from pilots—as well as advancements in the technology—meant that these policies soon required adjustment. Arizona released updated guidelines in October 2024 with three important changes:

1. Arizona defined the role of the newly created State Data and Analytics Office within the Department of Administration
2. Data governance was given a higher priority in administering AI solutions
3. State staff were given a clear understanding of responsibilities related to AI, including data protection, security, and privacy

Figure 6. Evaluating Generative AI Use-Cases

In Vermont's framework for acceptable uses of generative AI, "cite" refers to uses that are permissible, but require disclosure that work product was created with the assistance of AI.

Breadth of Distribution	Proofreading, Grammar	Brainstorming, First Draft, <25% AI	Collaborative Writing, About 50% AI	Human Edited, >75% AI	100% AI Content
Press release, prepared remarks	✓	cite	✗	✗	✗
Replies to public inquiry	✓	cite	✗	✗	✗
Public facing web content	✓	cite	cite	✗	✗
Memos, broad internal communications	✓	cite	cite	✗	✗
Internal process docs	✓	✓	cite	✗	✗
Source code	✓	✓	cite	✗	✗
Emails	✓	✓	cite	cite	✗
Chat	✓	✓	cite	cite	✗

Developing an AI strategy

Step 3: Raise Awareness Across Agencies

Responsible Use

Once guidelines are set, they need to be rolled out. State CIOs emphasized the importance of proactively communicating guidelines, and what they permit or don't, as well as making sure staff understand the AI team is available for advice. **"Rather than getting a lot of people [using ChatGPT] in the shadows,"** Josiah Raiche, Vermont's Chief Data and AI Officer, continued, **"people would reach out."** It is important that managers be empowered to make judgement calls rather than commit to overly restrictive guidelines. He tells managers to use the "VT Digger test," referencing the local newspaper. **"Would this be an interesting story? And if it'd be an interesting story, would it be a good or bad interesting story? Pretend you are an investigative journalist and got a call."**

Productive Use

After guidelines for responsible use are in place, the next step is to help staff fully understand the limitations and potential of AI's current capabilities. A common foundation helps set realistic expectations, dispel common myths about job displacement, and drive the cultural change necessary for effective AI adoption. Given the pace of change in the field, State CAIOs expect workforce education to be an ongoing project.

Build AI Enthusiasm

To get people excited, State CAIOs reiterated that online curriculum should be paired with live workshops. The demand is there: One small state had 3% of its workforce show up to an optional training. These workshops help to expand in-agency capacity, expanding the reach of COE teams. According to state AI lead, **"AI pilots tend not to be that technical. They tend to be prompt engineering and process engineering,"** meaning anyone in an agency can spearhead a new use case and become a champion of the technology. The key is motivation: "You are able to run more pilots if you get business owners excited." By raising awareness of AI's value, state IT leaders can get more staff and departments interested in joining AI efforts and amplify their impact.

Communicate AI's Limitations

At the same time, training also helps set realistic expectations. One CIO told us that departments participating in proof of concept projects using generative AI became **"really frustrated when they got responses that were not very accurate. They didn't really understand that they needed to invest in training the AI."** Training programs like New Jersey's can be used to help state workers understand both AI's potential and its limitations.

Case Study: New Jersey & InnovateUS Collaborate on AI Education

In an October 2023 executive order, New Jersey committed to creating a training program to help agency staff better understand and implement AI solutions. The State is a member of InnovateUS, a non-profit organization focused on innovation skills for public sector professionals. The State and the non-profit worked together to develop a curriculum and online training program to educate users in prompt engineering, low-risk use cases, and the importance of data protection. The educational materials and applications went live in July 2024. By March, more than 11,000 state workers had used the training and other states had begun adapting the curriculum for their own use. The state released the training alongside the NJ AI Assistant, a generative AI chatbot built for internal use and secured within the state environment. Over 20% of the state workforce has used the chatbot.

Today, New Jersey is running whole-agency live workshops to supplement the online materials. State IT leaders emphasized the importance of hands-on collaboration with trainees as they try actual use cases.

Step 4: Source Use Cases and Prioritize Them

Many states are developing processes to identify and prioritize use cases across agencies. Broadly, these efforts can be divided into "push" and "pull" approaches. In the former, COEs develop and "push" use cases out to the agencies; in the latter, COEs source ("pull") ideas from the agencies for projects of immediate benefit. Most states are pursuing a combination of both paths in parallel, and both approaches will likely be necessary to unlock the full impact of AI.

Push Use Cases

Across the states we interviewed, several standard use cases were widely perceived as low-risk and high-impact when "pushed" out to agencies. These initial use cases were not focused on automating service delivery, but rather helping agency staff be more effective. As one state AI lead shared, this early generation of AI solutions is already **"helping government employees quickly access and understand large repositories of data so that they can more efficiently help residents."**

Within the broad category of "pushed" solutions, state CAIOs have seen significant success in four categories that leverage the existing capability of AI technology and government competencies. These four areas represent "quick wins" where off-the-shelf solutions can be readily deployed to help a wide range of agencies, making them a powerful starting point.

Figure 7. Push Versus Pull Approaches

Pushing Standard Use Cases

IT departments in some States "push" use cases to their agencies in order to help them more efficiently support residents.

PRO	CON
• Potential to develop applications with wide applicability • Begins from understanding of AI's strengths and weaknesses • Easier roadmapping for resource-limited IT teams	• May not meet real demand • May be perceived as a "top down" solution & lack uptake

Pulling Ideas from Agencies

Some States provide a process for employees to propose use cases to be reviewed and assessed by IT departments.

PRO	CON
• Effectively demonstrates AI's utility • Creates advocates within agencies • Meets real demand if successful	• May lead to overly-bespoke application development • May spend limited resources on use cases ultimately unsuitable for AI

> "The call center agents could ask the chatbot the questions and provide the information... We're calling these things "policy bots" — basically an internal search function."
>
> – State AI leader

We also heard warnings about use cases that don't make sense today. One category of pilots respondents cautioned against were external facing chatbots or other AI solutions constituents interact with directly. As New York City's experience with the MyCity Chatbot showed, LLMs still hallucinate, and respondents generally regarded the risk of misleading users as too great to justify. Another category of use cases respondents warned against were agentic AI products used for screening people (e.g., deciding which resumes to advance.) The risk of algorithmic bias remains high, and most state CIOs felt the technology hasn't matured enough yet to be safely used.

Four "Push" Use Cases with Immediate Feasibility & Broad Applicability

1. **Content creation** of government or public-facing documents and resources (for review before publishing)
2. **Language translation and voice-to-text transliteration** of government or public-facing documents, resources, and speeches
3. **Coding and data analysis assistance** for IT staff responsible for updating legacy systems, government websites, etc.
4. **Internal "policy bots"** for agency staff responsible for answering questions about government procedures or policies, either to improve their own understanding or to provide answers to residents

Pull Use Cases

At the same time, many states also have a ground-up process through which individuals can propose use cases to the COE.

Beyond live workshops, most COE in large states provided a form for agency staff to submit ideas. This "pull" approach is an attractive way to make sure AI tools will actually be used by agency staff to overcome existing pain points.

Yet a "pull" approach can require a significant investment by the COE. As one interviewee noted, operating agencies are generally ill equipped to answer the question: "When is AI the right tool for a given problem?" COEs sourcing use cases must be willing to spend time and resources evaluating proposals.

Though each state has its own criteria for evaluating AI use cases, three categories predominated.

1. Is this a responsible use of AI?

Even setting data security aside, most states still err on the side of caution when it comes to implementing AI solutions at scale. For example, public-facing chatbots are an exception, not the norm.

2. Is this a productive use of AI?

Does the use of AI actually save time, or improve quality, compared to other potential solutions, technological or otherwise? One state, for example, required the inclusion of "business plans" in proposals to demonstrate an ROI.

3. Is this a productive use of the COE's limited resources?

Many COEs were inundated with requests, relative to their capacity to execute, forcing them to triage to the highest priority use cases.

Step 5: Launch Standard Off-the-Shelf Solutions

Deploy Off-the-Shelf Solutions

The fastest way to push AI solutions into the hands of users is to work with commercial off-the-shelf solutions. Think, for example, of the enterprise versions of consumer-facing apps like Anthropic's Claude, Microsoft's CoPilot or Midjourney. In contrast to bespoke AI solutions for highly particular problems, these products can help across state agencies with tasks like brainstorming, editing, drafting, coding, and organizing information.

Create Sandbox Environments for Piloting

As a first step, many states created a "sandbox environment," a secure, isolated virtual space where AI solutions can be tested without affecting live IT infrastructure. There are two types of sandbox generally used for commercial products. Most common is a user sandbox, where a subset of state workers use AI tools in the sandbox before they are scaled to the broader organization. Some states have also created policy sandbox environments for tools that may require exemption from the state's IT policy, allowing the AI team to test out tools' capacity in safety. As AI becomes increasingly ubiquitous, some states are skipping the user sandbox for commercial off-the-shelf solutions entirely.

Decrease Time, Improve Quality

Based on early pilots, states are following the commercial sector in identifying AI tools that can both save time and improve quality. Pennsylvania, for example, found that users reported saving more than an hour and half a day after piloting ChatGPT Enterprise, while Vermont improved code quality and documentation in Salesforce assisted by AI.

Case Study: Pennsylvania Pilots ChatGPT

Pennsylvania announced a pilot of ChatGPT Enterprise in January 2024. During the next year, 175 staff from 14 agencies used the tool to test AI's impact, about half of whom had never used ChatGPT. Data was collected biweekly through feedback sessions, user interviews, and live demonstrations.

Through feedback, staff self-reported saving an average of 95 minutes per day on administrative tasks, such as drafting emails, summarizing lengthy documents, and navigating complex bureaucratic processes. Over 85 percent of users reported a "somewhat positive" or "very positive" experience, demonstrating broad staff buy-in across various government roles. Following the pilot, Pennsylvania's Office of Administration is working to acquire a vetted suite of generative AI tools.

Case Study: Vermont Automates Code Quality Checks

Vermont's IT team spends significant effort enforcing bespoke security rules on code in the State's Salesforce environment. The State decided to pilot an AI solution that could automate code review, potentially saving time and getting to a higher standard of compliance. The tool fed the results of the AI review back into the coding environment as comments for the developers. As Josiah Raiche, the State CAIO told us, creating the AI tool required a significant investment of time and effort: "We spent time and wrote really good descriptions of the rules and what a violation looked like." However, the results were impressive: 70-80% resolution rates, and a decrease in the amount of "technical debt that would accumulate with time."

Step 6: Prototype Custom Point Solutions

Buy or Build—At a Low Cost

When it comes time to move beyond off-the-shelf solutions to addressing bespoke use cases through prototypes, states must decide between building and buying. In many states, limited AI staffing and budgeting means that CAIOs have built prototypes using volunteer workforces (such as students or non-profit assistance). Buying, on the other hand, generally requires creating a structured procurement process. It is preferable to use a "challenge-based" procurement model that defines the problem the state is trying to solve, and allows vendors to propose creative fixes. As in the example of California below, states can use the prospect of larger contracts down the road to incentivize vendors to experiment at a low cost.

Case Study: New Jersey's Collaboration with USDR (Building)

U.S. Digital Response (USDR) is a non-profit organization with nearly 11,000 volunteers that work alongside governments to augment their digital capacity. New Jersey's Department of Labor partnered with USDR to improve language access for unemployment insurance (UI). The State had learned that residents with limited English proficiency struggled to access UI benefits. USDR helped the State launch a Spanish-language UI application using AI, which reduced the average application time from 3+ hours to just 28 minutes. New Jersey now reports parity in form completion time between English- and Spanish-speaking applicants.

Case Study: California's RFI2 Procurement Process for Gen AI (Buying)

California identified five priority use cases (e.g., improved call center productivity and traffic mobility insights) and then released a Request for Innovative Ideas (RFI2) to identify vendors to design AI solutions and develop proofs of concept. Vendors were offered only one dollar for the proof of concept phase, but participation led to their consideration for a larger contract if the solution was scaled.

State IT leaders pointed to the California Department of Tax and Fee Administration's RFI2 for Call Center Team Productivity as a particular success. The Department handles more than 600,000 calls a year (alongside emails and chats), and representatives often have to spend a long time searching through tax information scattered across articles, guides and manuals to answer constituent questions. The RFI2 requested AI prototypes that could help staff answer questions, improve monitoring, and generally expand the Department's capacity to help constituents. State IT leaders see the pilot as providing an initial foundation for providing help to call center staff across the state more broadly.

Learning from Failures

Even in a more intensely resourced environment, most respondents still recommended starting with low cost projects due to the high failure rate of AI pilots.

> "There is such a high failure rate on AI pilots that you definitely want to minimize your cost and start small. You need to carefully measure your pilots and see if they actually increase efficiency... We did about 10-15 pilots, with a 50% success rate."
>
> – State AI leader

While we heard about a range of generative AI success stories detailed in case studies on page 34, state IT leaders emphasized the importance of going into these initial efforts clear-eyed. Many AI pilots and prototypes will not succeed. Among our respondents, one state tried to use AI to automate its wetland permit application process before finding it didn't actually save time. Another state did a pilot project with law enforcement to automate portions of officer write-ups, but found the underlying data wasn't organized enough to get effective results. "At the time," the respondent told us, "it was hard for the generative model to distinguish between whether a firearm was pulled during the incident or found in the car after." A third state tried to use an agentic AI bot to detect and send cease and desist letters to websites illegally selling liquor, before finding the technology wasn't advanced enough yet. However, the state holds out hope that AI may be able to help in the near future, as agentic AI is improving. In general, this example illustrates that, rather than reject pilots as "failures," it's important to see them as learning experiences that may be picked up in a few years as the technology matures.

> "In general, failure today may be success tomorrow given how rapidly the technology is developing, so keep use cases in the holding pen."
>
> – State AI leader

Overall, we found that custom AI solutions have the best chance of succeeding in state agencies where the relevant data is already organized. As one state AI leader said, **"Data is really going to be the big thing moving forward...For those departments that it took a little bit of time to do even just some basic tagging, they were able to get way more out of the AI tools."** This substantiates the notion that effective data management is a precursor to building out AI applications successfully.

A Wide Range of Use-Cases

Beyond the standard suite of "push" use cases, we found many states were already deploying generative AI to solve a broad suite of sophisticated problems. Here are a few examples:

Case Study: Connecticut Prevents Fraudulent Unemployment Claims

States often receive a significant volume of fraudulent unemployment claims, and Connecticut is no exception. Its Department of Labor piloted a suite of AI and ML tools from LexisNexis in an attempt to decrease fraud. By integrating the tool into the unemployment insurance application, they were able to verify applicants' identities and detect fraudulent claims before they were submitted. Once scaled, the State expects this solution to save millions of dollars.

Case Study: Arizona Reduces Bias Through Body Camera Footage

Arizona's Department of Public Safety launched a partnership with TRULEO and Arizona State University in October 2024, leveraging AI to conduct a field study on body camera footage. TRULEO converts body camera video to text. After automatically redacting police officers' and residents' personally identifiable data, the tool detects and documents interactions like introductions, explanations, and de-escalation attempts. This AI solution is intended to expand the agency's capacity to review footage and reduce biases.

Case Study: Georgia Ensures Performance of a Critical Website

Every month, Georgia's primary resident-facing website, Georgia.Gov, receives six to eleven million denial of service attacks, or attempts to inundate the website with excess traffic. A recent pilot used AI in the background to identify traffic intent. Visitors from flagged IP addresses were posed with challenges, and after proving that they were humans, they were able to proceed. This AI solution ensures continued performance of this critical website so that residents can continue accessing the information and services they need.

Evaluating Impact

Time-savings are a straightforward way to measure the efficacy of some AI pilots. But CAIOs consistently told us that other desired outcomes are harder to put a number on. For example, many pilots have the goal of finding a solution that works "as well" as a human. But even setting aside the difficulty of measuring that outcome, some state IT leaders wonder if that's the right standard. Similarly, we were told that measures of bias are still inadequate.

> "Measuring the impact of generative AI pilots is critical and there isn't robust science yet on how to do this."
>
> – State AI leader

For example, one state tried to use AI to predict when bridges would require repairs. (Presently, these projections are made by infrastructure experts.) While the pilot "worked just as well as the best system that is currently available," the decision-making process was harder to review, and the filing function behaved strangely, making the project hard to classify as a success or failure. Most CAIOs agree on the importance of taking the time to outline clear methods of evaluation before embarking on a custom build.

From Pilot to Scale: The Role of Talent in AI Deployment

The best practices and expertise developed during the first six steps lay the groundwork for deploying AI. But to deliver AI at a government scale, most CAIOs have concluded that they need to create additional, AI-enabling roles throughout their workforce, detailed below.

The Three Foundational Roles

Due to a lack of budget, states have generally begun their AI journeys by drawing on existing talent and appending AI to titles and job functions. Most CAIOs argue three IT roles are necessary to carry a state through steps one through six: an AI director (or equivalent), a product manager, and a policy manager. Alongside these three foundational roles, state IT leaders highlighted the importance of owners in operating agencies. This helps expand the impact of a resource-constrained central AI team, and creates buy-in for change among agencies.

Past experience in AI is helpful for these roles, but not required. Respondents told us that staff with strong business administration and/or product management skills can be successful too, as long as they have a broad understanding of the IT function.

Figure 8. The Three Foundational AI Roles

Candidate Profiles	The Director	The Policy Manager (Responsible Use)	The Product Manager (Productive Use)
Previous Experience	6+ years of experience with AI, data science, IT, and/or IT management	3-5 years of experience with AI, data science, and/or IT governance	3-5 years of experience with AI, business administration, and/or product management
Anticipated Role	Responsible for developing and implementing the AI strategy. This will involve (1) building awareness of AI among elected officials, state employees, and other stakeholders, (2) identifying and prioritizing use cases, and (3) overseeing the ongoing monitoring of use cases.	Responsible for developing and implementing standard governance processes for the responsible use of AI. This will involve (1) determining how the State should review and assess potential use cases and (2) ensuring ongoing compliance and adherence to the state's AI guidelines.	Responsible for assessing and managing potential use cases for the productive use of AI. This will involve (1) developing business cases and making build-versus-buy decisions and (2) managing the implementation of use cases.
Reporting Structure	Reports to the Chief Information Officer	Reports to the Director	Reports to the Director

Figure 9. Digital Services Roles to Enable AI Solutions

	Solutions Architect	Business Analyst	User Researcher	Technology Trainer
Anticipated Role	Works closely with employees to understand their challenges and then designs AI solutions to address those needs.	Gathers requirements for specific AI solutions and develops business cases to support related investments.	Codifies which AI solutions are and are not working effectively so that learnings can be incorporated into future use cases.	Helps state employees understand what is and is not feasible with AI, and teaches them how to incorporate AI into their work.

Hiring for Scale

More broadly, to truly implement AI requires enabling functions like data science—and some states are further ahead than others. According to one CAIO, **"I'm highly cognizant of not messing up early AI, and I've been working the past seven years on statewide data governance...I've had to add data scientists and data engineers and all these other job classifications to the state."** Other states may need to play catch up, creating new data science roles or amplifying existing efforts.

The same is true for a number of other "digital services" roles, which—while they may not be exclusively AI focused—are necessary precursors for a large-scale roll out of many AI applications.

Doubling Down on Productive Uses

Deploying AI at scale will require a mindset shift in the legislature, executive chamber, agencies, and among job seekers. As one respondent told us: **"More people imagine a career...in regulation rather than in deployment. I would love for there to be more homes for people to be the carrot, rather than stick."** While oversight remains critical, AI presents an opportunity for state governments to develop a solution-minded, experimentation-first culture. As shown in the case studies, COEs can validate this approach—and potentially attract both political support and new talent—by improving quality and achieving efficiency over painful manual processes.

> "I tell agency staff to think about AI more like a 'power tool.' You need some level of training by an operator, and you use it for a cluster of similar tasks. But then it helps you zoom out of the mechanics of doing the work, and allows you to get things done quickly and well."
>
> – State AI leader

Prepping the Workforce of the Future: Early Career Talent

Demand for Early Career Talent

Once foundational AI roles are filled, our study strongly suggested a significant role for early career talent in helping promote the proliferation of data science, digital services, and ultimately AI, throughout state government.

As one state AI leader put it: **"Because AI is a new topic, early talent with technical backgrounds could still add value on teams because no one has much experience with AI."** State leaders still emphasized that the core AI leadership roles likely require staff experienced in management and large organizations. However, "AI-Native" early career talent is highly valued for spreading the technology more widely at the agency level. State CAIOs are searching for ways to strengthen their relationship to local universities (through internships, job fairs, and nonprofit matching programs, as described below), a key funnel for AI-native talent into junior government technology jobs.

> "I keep two cubes available for interns at all times. We do a big summer internship program. Some of our best employees have started as interns and liked it, and kind of stayed on as extra help until they graduated, and then transitioned right in."
>
> – State AI leader

Supply of Early Career Talent

Today, there is an *excess* supply of students interested in working for the government, relative to roles currently available. Survey data reveals that the percentage of government job applications from recent graduates (relative to overall job applications) has been steadily increasing.

Figure 10. College Graduates Show Increased Interest in Government

Year	Government job applications as a percentage of total job applications from graduates
2024	**7.4%***
2023	**5.5%**
2022	**<5%**

*Note: based on ~2700 surveyed students

The resurgent interest in government jobs among young people is likely driven by several compounding factors. A top concern among recent graduates is ending up in a career which they are not passionate about, a fear ranked second only to concerns over work/life balance. Government gives people a chance to work directly on issues that impact their community. Conversely, the class of 2024 ranked "job stability" as the single most attractive factor in a job, and government jobs have historically been stable at the state level.

Among technology talent specifically, there are further, market-based factors at play. Job availability at technology companies dropped by 30% from 2023 to 2024. In addition, the number of students graduating from academic programs that prepare students specifically for technology roles in government is increasing. Through the Public Interest Technology University Network, approximately 60 universities and colleges have committed to the development of academic curricula, research agendas, and other learning programs related to public interest technology, a number that has grown by nearly 50 percent within the past three years.

There is an ample supply of AI-native talent eager to work on impactful problems. The challenge is that government is not always prepared to take on early career talent—and the matchmaking process between supply and demand is not occurring effectively.

Constraints on Early Career Talent Entering Government

On the government side, state governments pointed to two primary constraints impacting their ability to hire early career talent for AI and AI-enabling roles. The first was mentorship capacity. Hiring recent graduates means senior IT staff need to spend time providing coaching and training, impacting their own already-strained capacity. One interviewee noted that increased hiring of early career talent would also likely require hiring more senior staff engineers. The second concern was the rapid onboarding and upskilling necessary for early hires to make an impact on quick moving AI initiatives. Historically, one interviewee told us, it can take up to five years for new IT staff to begin fully contributing. To effectively harness AI technology and make use of talent coming out of universities, states will need thoughtful hiring and training practices that emphasize immediate impact.

University career advisors we interviewed pointed to significant constraints on the applicants' side as well. State government recruiting processes can often seem opaque compared to private sector job openings. One respondent told us: **"The process is opaque; the websites are outdated, the timelines are long, and [students] have trouble finding contacts."** In general, the career advisors reiterated that students are less comfortable finding and applying to a government job that might match their interests and skills.

Taken together, these facts present significant barriers to states locating and hiring young people eager to work on government AI projects, and act as AI-native ambassadors across agencies.

Constraints Expressed by Governments

The research revealed two primary government constraints on hiring early-career talent.

1 Mentorship Capacity

Senior employees need to coach and train new hires, which impacts their own capacity.

> **"There is definitely a need for more junior talent in big state agencies like the Department of Labor, but still there is the question of who provides the requisite mentorship. [We] need more senior engineers to act as mentors."**

> **"Given everything else that's already on the state employees' plates, if you're also asking or expecting them to dedicate X amount of time to coaching and training, then it's less appealing."**

1 Onboarding Timeline

The lean nature of states' AI teams means that new hires need to work independently quickly.

> **"We need to change the model to onboard for 1-2 years to get talent up to speed quickly, so they can learn and deliver. That's not something we have done historically. Usually we think people need to work here for 5 years to be successful."**

> **"If you can develop a program where you can be sure that the early-career talent is value additive, and it's reducing workload for others and not adding to it. Then it obviously makes sense."**

Constraints Expressed by Universities

University career advisors summarized four issues in government recruitment of early-career talent.

1 Recruiting Timeline

Private-sector companies recruit at the start of year for jobs that begin post-graduation; governments recruit later, when many students have already accepted other offers.

2 Role Description

Government job postings are often unclear about what the roles entail, or they use language that differs significantly from "similar" private-sector postings.

3 Application Process

Many government roles are posted only on government websites and/or job boards, rather than the recruiting platforms with which students are most familiar (e.g., Handshake).

4 Perceived Barriers

Students are less familiar with how to network for government jobs vs. private-sector jobs. They may not know whom to speak with, and government recruiters often have a limited presence.

The Role of Matchmaking

Out of the universities we surveyed, all had the federal government represented at their career fairs–but only four had relationships to state or local governments, a significant missed opportunity. As one interviewee put it: **"CIA, DOD, FBI, and other federal agencies recruit on campus and get a really good return; they find great candidates who often accept offers."** While agencies like the CIA make ample use of on-campus events such as coffee chats and simulation exercises, state governments are often totally absent from on-campus recruiting.

More positively, states and universities that have set up structured recruiting programs have had rapid success in creating a pipeline from campus into government tech roles, including in AI. Massachusetts, for example, worked with Northeastern University on an AI higher education program to great effect.

Case Study: Massachusetts Builds and Recruits with Northeastern's Help

The Commonwealth of Massachusetts worked with Northeastern University throughout 2024 to mentor 15 students from various fields of study. The students were tasked with building a chatbot for RIDE, the MBTA's paratransit service. Initially, the chatbot was supposed to provide residents information about available RIDE services. But after user research, the students found that residents didn't want a technological solution like a chatbot. They wanted to talk to a person, but were frustrated because they often ended up being transferred when an agent didn't have an answer to their question. The students pivoted to building a chatbot to help call center agents quickly locate relevant information, which is now moving into production.

According to William Cole, the State CTO, the program was also successful in driving recruitment, with several students from the first cohort returning for internships or full-time roles with the state, including 3-4 to the AI team. The partnership was so successful that Massachusetts has now expanded the initiative to a partnership with the University of Massachusetts at Amherst.

As InnovateMA's example shows, even a small initial investment in building a pipeline between universities and states can have a large impact in building an attractive brand and raising awareness among early career talent. The incentives are there. As one interviewee observed, **"Direct relationships between universities and state or local governments are effective and advantageous for all stakeholders."** Similar matchmaking programs are also helpful for more senior roles, and existing federal models like Tech2Gov, featured below, provide templates that can be adapted locally.

Case Study: Tech2Gov Hiring Forums Bring Technologists into Government

Between January 2022 and October 2023, the Tech Talent Project in partnership with The Volcker Alliance held a series of "Tech to Gov" virtual hiring forums to match candidates with government technology jobs. These events have connected over 10,000 technologists with 130+ Local, State and Federal agency talent and hiring teams. With more than 200 technologists hired into agency roles, the events served as a training opportunity for agencies to adopt modern hiring practices for technical talent. On the strength of these early results, Tech Talent Project and The Volcker Alliance have now, more formally, brought this approach to federal agencies.

Looking Ahead: Filling the Talent Gap

There are a series of actions states can take—possibly in conjunction with private philanthropy—to accelerate the roll out of responsible and productive uses of AI across core state government functions, but doing so requires a thoughtful approach to talent. Talent needs vary primarily based on the digital maturity of state IT departments, and where specifically each state is in their AI journey.

Beginners

States that do not yet have a point person for AI within their IT departments will likely benefit from hiring one or assigning responsibility to an existing staff member, to create a holistic approach to the state's AI strategy. While AI is not fundamentally different from other emerging technologies, given its rapid pace of change, there is a benefit to assigning someone at the very least to help ensure the state's policies stay up to date, and to learn from other states' progress. Ideally, that individual would help guide the state through the six steps outlined in this report, and would tap into the wealth of philanthropic resources currently available to promote technology modernization in government, including the adoption of AI. Accessing those resources generally requires a clear responsible party—yet most states still lack this role today.

Case Study: The Center for Public Sector AI's Advisory Peer Group

The Center for Public Sector AI (CPSAI) has created a peer group in which states with designated AI advisors meet on a bi-weekly basis, to share resources and discuss progress they are making in their respective AI journeys. The peer group has created a forum for structured conversations around emerging policy topics, like how to evaluate the success of an AI implementation, and for distributing learnings about AI applications—both those that have succeeded, and those for which the technology is not yet sufficiently mature. The CPSAI peer group was consulted throughout this study.

Early Adopters

Most of the states that participated in our study had already progressed through some iteration of the six steps described in this report. They are just now reaching the stage where they are attempting to scale their staffing.

While state AI leaders have submitted budget requests, generally for 3-5 person teams, many anticipated multiple budget cycles before the roles will be granted. In the interim, there is clearly space for a strategically targeted, philanthropic-funded fellowship program to help AI leaders in early adopter states expand their capacity by hiring a deputy. As one such leader put it, **"I need someone who can plug AI into all the things that AI needs to get plugged into. They don't need to be super technical, they need to figure out the business case, the process evolution, and the change management component."** Another commented, **"There is so much that is happening in other states. I haven't had the bandwidth to look across what states are doing and figure out how to learn from it. Someone to help me with that would be great."** While philanthropic funding will not

solve a state's long-term staffing needs, there is a clear role for philanthropy to help leading states "bridge" into long-term staffing.

Case Study: Pennsylvania Leverages Collaboration

Harrison MacRae, a former Govern For America fellow who is currently the Director of Emerging Technologies for the Commonwealth of Pennsylvania, oversees the state's Gen AI strategy, administers the Generative AI Governing Board, and led a year-long pilot of ChatGPT Enterprise with state employees. Throughout Harrison's work with the Commonwealth, he has collaborated with Carnegie Mellon University's faculty and student capstone groups to engage AI expertise. Members of his team have also participated in public impact fellowship programs such as New America's Share the Mic in Cyber fellowship. This proactive engagement has aided in recruiting talent to the Emerging Technologies team to expand the state's AI capacity, and building awareness of the Commonwealth as an employer focused on innovation and employee-centered approaches to AI adoption.

All States

Lastly, all states could benefit from more structured programs to bring more early career IT talent into government, creating the next generation of public servants and evangelists for technological change. There is excess interest in government positions among college graduates relative to the available roles. States that do have structured matchmaking programs, like Massachusetts, have seen how early career talent can make a significant impact, particularly in fields like AI where "no one has much experience." Govern For America's own programs for placing and training early career technical talent in state and city governments could provide a model state IT leaders can draw upon as they build out AI capacities.

Govern For America's own programs for placing and training early career technical talent in government validate this finding—and create an infrastructure state IT leaders can draw upon as they build out their AI capabilities. We are excited to keep serving states across the country as they explore how to harness AI's potential to improve government efficacy.

GOVERN FOR AMERICA

Get in Touch. Support our Work.
Join the Movement.

www.govforamerica.org