

Beyond Generation:

The Rise of Agentic AI in State Government

March 2026



NAS CIO[®]
Representing Chief Information
Officers of the States

Beyond Generation: The Rise of Agentic AI in State Government

For more than three years now, state governments have been operating in a world with readily accessible generative AI (GenAI) tools. What started as conversations around possibilities in early 2023, has morphed into the **number one priority** for state chief information officers (CIOs) in 2026.

Today, states are heavily involved in artificial intelligence (AI) governance. We know from NASCIO's **2025 State CIO Survey** that 88 percent have AI responsible use policies, flexible guardrails, security policies and/or ethics requirements. We know that 84 percent of states are inventorying and documenting uses in agencies or applications. We also know that 82 percent have created advisory committees and/or task forces. Given that these statistics are from last year, we can assume these numbers are actually higher today.

State employees are also becoming more comfortable using GenAI tools in their daily work. State CIOs reported in 2025 that 82 percent of employees in the CIO's organization are using GenAI tools in their daily work (up from only 53 percent in 2024).

Nearly all states have initiated GenAI pilot projects (90 percent in 2025) and 86 percent reported implementing proofs of concept. States are also starting to scale GenAI tools beyond the pilot phase. From a **2025 joint report** NASCIO published with Accenture, *Harnessing GenAI to Elevate the Citizen Experience*, we know that 23 percent of state CIOs reported using GenAI virtual agents and chatbots at scale (again a number that is likely higher today).

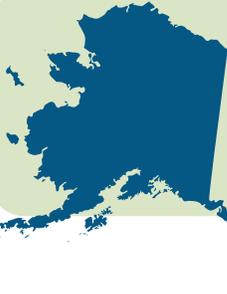
While large language models have been revolutionary at document creation, summarization, transcription, translation, data analytics and software code generation, an evolution is happening now. As generative AI chatbots and other tools slowly become more autonomous and able to do more tasks, states are finding that they are utilizing agentic AI—AI that can automate complex tasks. We know from data collected during our research process for the NASCIO Top Ten list that eight states reported that they already have some agentic AI tools in production.

The following table is a useful comparison between GenAI and agentic AI:

Characteristic	Generative AI	Agentic AI
Core Capability	Generating text, images, code or music based on learned patterns	Planning, decision-making, multi-step execution without human intervention
Memory & Context	Limited memory (short-term context retention)	Persistent memory (remembers past interactions, adjusts plans accordingly)
Autonomy Level	Requires human prompts to generate responses	Operates with minimal human input, executing complex workflows
Integration with External Systems	Minimal integration (relies on APIs or tools for external functions)	Deep integration (connects with APIs, databases, physical systems)
Learning Ability	Static—learns only through retraining by developers	Evolves—learns from interactions and refines behaviors
Typical Use Cases	Content creation, summarization, coding assistance, brainstorming	Workflow automation, personal assistants, business operations
Business Impact	Enhances efficiency in content-heavy tasks but does not automate workflows: <ul style="list-style-type: none"> • Average increase speed: 25% faster • Average quality improvement: 40% (source) 	Drives automation, reduces human workload, enhances business scalability: <ul style="list-style-type: none"> • Time savings: 30-60% • Process acceleration: 40-90% faster (Bornet et al.)
Examples	ChatGPT, Claude, Gemini, DALL-E, Midjourney, Copilot	AutoGen, MS Copilot Agent Builder, UiPath Agent Builder, OpenAI Operator, Google Vertex, Crew.ai, Relevance.ai, Agentforce

Source: Adapted from *Agentic Artificial Intelligence: Harnessing AI Agents to Reinvent Business, Work, and Life*, by Pascal Bornet et al. (Boston: Harvard Business Review Press, 2024)

States are likely to find plenty of business cases for these advanced AI tools. A [recent study from Accenture](#) found that despite recent leaps in digital government and generative AI usage, 49 percent of citizens still have trouble navigating citizen websites. And for government workers, 62 percent said that “inefficient processes have a significant to severe impact on their ability to deliver core services to users.”



[Alaska](#) is considering adding agentic AI to their myAlaska portal. The state is considering features such as “autonomous completion of multi-step government service transactions on behalf of users” and “dynamic form filling, document retrieval and eligibility checks.”

A Gradual Evolution

State technology leaders are unlikely to walk into the office one day and say, “This is the day we start using agentic AI.” It’s more likely that GenAI tools are starting to add capabilities, extending existing GenAI tools incrementally while keeping humans in the loop. Eventually states will have crossed into “agentic” territory. In some cases, perhaps the tool itself hasn’t changed, but the responsibility the tool has been allowed to take on has.



[Tennessee](#) is looking for a next generation ERP solution with agentic AI capabilities. They want a solution that can “detect anomalies in payroll, flag procurement bottlenecks, surface compliance risks or even identify potential fraud.”

Because this shift won’t happen overnight, states are likely to pass through several phases of maturity from using GenAI to using agentic AI. Many states are already several steps into this shift. The following is an overview of five phases that states may go through as they build upon generative AI and increase agentic AI capabilities:

Phase 1: Assistive generative AI

AI helps humans complete individual tasks faster. Most common uses include:

- Drafting documents, correspondence and other content
- Summarizing policies and legislation
- Answering frequently asked questions for employees or citizens
- Writing computer code

In this phase humans are the primary doers and decision makers. The AI is creating but humans are taking action. Most states have reported using AI in these ways over the last few years.

Phase 2: Context-aware generative AI

AI starts remembering context within a task. This is where many states are starting to expand pilots. Examples include AI's ability to:

- Maintain conversation context (citizen bots that remember what a user already answered)
- Use agency-specific knowledge bases (internal or external chatbots that are based on agency policy or agency websites)
- Produce outputs tailored to roles or programs

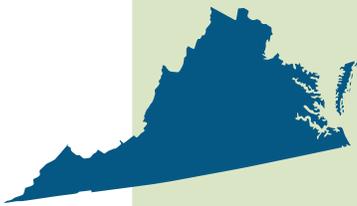
In this phase humans are still the primary operators, however the AI is becoming more sophisticated.

Phase 3: Task-level automation

AI is allowed to take limited actions inside defined boundaries. This is the subtle tipping point into agentic AI. AI is no longer only suggesting but starts doing (narrowly). Examples include:

- Filing a help ticket after a conversation
- Pre-filling a form using collected data
- Routing a document for approval
- Updating a knowledge base after human sign-off

In this phase humans are the reviewer or approver.



In July of 2025, Virginia's then Governor Glenn Youngkin [issued](#) an executive order to use agentic AI to improve government efficiency. The pilot would use an agentic AI tool to quickly scan existing regulations and guidance to spot conflicts with statute, cut redundancies and flag opportunities to simplify and clarify the language."

Phase 4: Stateful, multi-step workflows

This is fully agentic AI. "Stateful" AI agents retain context about interactions with users—not just within a specific conversation. AI manages progress across time, steps and systems. At this point AI tracks what's already done and what comes next, coordinates across systems, handles exceptions and escalates to humans when rules require it. Examples include:

- Employee onboarding managed end-to-end
- Compliance updates monitored, routed, tracked and logged
- Citizen applications guided from intake to decision

In this phase humans provide oversight, exceptions and final authority.

Phase 5: Proactive and adaptive agents

AI initiates work before anyone asks. This phase will come last as it raises the most governance and risk questions. Examples in this phase include:

- Detecting backlog risks and reallocating work
- Flagging policy changes before deadlines are missed
- Reaching out to citizens when action is needed

In some cases, agentic AI will develop as existing tools gain memory, permission to act, responsibility for outcomes and the trust of the humans utilizing it. In other cases, it will arrive as a new tool with new capabilities from a trusted vendor.

Challenges for States in the Evolution to Agentic AI

While states have started scaling GenAI relatively fast, giving AI agents autonomy is something that may need to happen more slowly. States are faced with the realities of governance, risk, legacy systems and workforce resistance to change.

Like with non-autonomous GenAI, states will likely try giving agentic AI bots more autonomy internally first before applying them externally and for citizen services. It will be even more important to have governance and guardrails in place with a heavy human hand in the early days.



In July 2025, [Delaware](#) launched an AI sandbox initiative to test and encourage the responsible development of agentic AI.

State governments also must deal with the realities of legacy systems. Agentic AI depends on applications having the ability to interface with one another, workflow engines and integration of data. These are all things that states are improving incrementally right now.

And of course, states can only scale at the speed of trust. In recent research from Accenture which was included in our joint report, only 55 percent of surveyed frontline government employees agreed with the statement, “I feel positive about the potential integration of AI into my day-to-day work.” In the same report, we learned that 75 percent of state CIOs have serious concerns about using GenAI in direct citizen services, with top concerns being AI accuracy, data security and privacy and availability of adequate training and support. These concerns only become more relevant when AI gains autonomy.



The [Texas Department of Transportation](#) (TxDOT) is piloting a project that integrates an AI-enabled agent capability with TxDOT’s TRACK complaint management system to explore improved case assignment, workload balancing and issue resolution. The goal is to enhance the speed and quality of the complaint-resolution experience for citizens.

Security Risks

AI agents introduce new security risks for states because they can be manipulated into taking unauthorized actions. These actions might include accessing sensitive data, misusing trusted systems or escalating privileges, and before a human realizes it, the damage is done. Because these agents can act autonomously and across multiple systems, a single compromise or bad instruction can quickly spread, leading to large-scale data breaches, fraud or service disruptions.

To reduce cybersecurity risks when using agentic AI, states should keep these systems tightly controlled, closely supervised and easy to stop if something goes wrong. Agentic AI should be treated as powerful but untrusted by default, with clear limits on what it can access, what actions it can take and how long it can operate without oversight. Strong guardrails, human checkpoints and continuous monitoring help ensure AI agents remain helpful tools rather than autonomous actors that introduce security, privacy or operational risk.

The following recommended mitigation practices are adopted from the [OWASP Top 10 for Agentic Applications for 2026](#):

- **Limit access by design:** Give AI agents only the minimum permissions they need for each tool or task, such as read-only access or narrowly scoped application programming interfaces (APIs).
- **Require human approval for high-risk actions:** Ensure that humans must approve destructive or sensitive actions like deleting data, transferring funds or publishing content.
- **Use secure execution environments:** Run AI tools and code in isolated sandboxes with strict controls on where data can be sent.
- **Validate intent before execution:** Check AI-generated plans and tool requests against security and policy rules before allowing them to run.
- **Apply usage and cost limits:** Set ceilings on how often tools can be used or how much they can spend, with automatic throttling or shutdowns if limits are exceeded.
- **Use temporary, session-based access:** Issue short-lived credentials that expire quickly and are tied to a specific user or session.
- **Verify tools and meaning, not just syntax:** Ensure the AI is using the correct, approved tools for the intended purpose and fail safely when requests are ambiguous.
- **Log and monitor continuously:** Keep detailed, tamper-proof records of AI actions and watch for unusual behavior or unexpected changes over time.

Opportunities for Improved Outcomes

Inevitably, the AI tools that states are using and scaling today will become more advanced, more autonomous and more useful. With early iterations of GenAI, the AI got to be the “creative” leaving all the “doing” up to the humans. With agentic AI, AI can start doing more of the tasks, leaving creative, human-centered and face-to-face work for humans. The productivity gains can be significant. As one state CIO said in our joint report on scaling GenAI for citizen services, “I like to look at where we are the most behind. In two years, I would like to see that we’ve improved productivity by 10x using agentic.”

Given the workforce shortages in state government around the retirements of the baby boomer generation, and the demands put on government as our population ages, augmenting state government work with agentic AI can be a real solution to working through backlogs and getting services in the hands of citizens faster. As another state CIO said in the NASCIO-Accenture report, “(Agentic AI is) not replacing you; it’s just making you 10 times more productive.”

Patterns of Success

The recommendations found in NASCIO’s previous reports on GenAI; [*Generating Opportunity*](#) and [*Your AI Blueprint*](#), still stand when it comes to agentic AI:

1. States should continue to **set policies** around AI and ensure that they are updated to include language around the risks and rewards of agentic AI as this new functionality becomes increasingly common.
2. If **governance around AI** has not yet been established, it has never been more important. Guardrails and privacy considerations need to be in place as AI becomes more autonomous. Regular testing and oversight continue to remain essential.
3. **Data quality** has become a top priority for state CIOs in recent years. Data quality will continue to grow in importance as AI agents rely on data accuracy and data sharing to complete workflows.
4. **Build partnerships** not only between state CIOs, chief information security officers, chief privacy officers and chief data officers, but with agencies or offices that must be involved in developing the processes and workflows that will be carried out by agentic AI in the future.
5. Adopting an **incremental approach** is perhaps never more critical than it is with AI that will begin acting autonomously. Taking a slow and incremental approach allows for time to learn lessons and build trust.
6. **Continue to be transparent.** Governments must prioritize trust and transparency in design and delivery of AI tools, especially agentic AI. Communicate openly about AI usage with both government workers and citizens.
7. It’s never been more important to **add security controls** to ensure the agents are controlled, supervised and can be stopped.

As you have read through this report, and the different phases of evolving toward agentic AI, perhaps you have recalled some of your own experiences with AI and realized that you have already benefited from some early uses of agentic AI either in the private sector, using AI tools at work or at home or in interacting with a government. These changes are happening now. For state governments, it’s a great time to step back and review policies and guardrails to make sure your organization is ready for agentic AI.

Primary Author
Amy Glasscock, CIPM
Program Director,
Innovation and Emerging Issues
NASCIO

aglasscock@NASCIO.org



About NASCIO

Founded in 1969, the National Association of State Chief Information Officers (NASCIIO) represents state chief information officers (CIOs) and technology executives and managers from the states, territories and District of Columbia. NASCIIO's mission is to advance government excellence through trusted collaboration, partnerships and technology leadership. NASCIIO provides state CIOs and state members with products and services designed to support the challenging role of the state CIO, stimulate the exchange of information and promote the adoption of IT best practices and innovations. From national conferences to peer networking, research and publications, briefings and government affairs, NASCIIO is the premier network and resource for state CIOs.