

CENTER FOR CIVIC FUTURES

THE STATE OF STATE AI 2026

A field report from the people running AI inside American state and territorial governments

August 2026

CENTER FOR CIVIC FUTURES

Executive Summary

Center for Civic Futures (CCF) works to turn isolated decision makers into a connected field, one equipped to deploy emerging technologies that improve service delivery while preserving human agency and dignity. This report is part of that work. It gathers what state leaders told us they are facing and what they need. Named plainly and in one place, those challenges stop being individual burdens and become shared work. States learn from each other, and partners can see where support belongs.

In CCF's interactions with our government partners, we've found that over the past two years, state, territorial, and tribal governments have increasingly moved beyond questioning if they should adopt artificial intelligence (AI) and onto figuring out how to responsibly put it into practice.

They are focused on building governance structures to guide decision making, upskilling staff, and actively experimenting with use cases. As the technology advances and governments progress along their maturity curve, the issues they are grappling with grow more complex. Many agencies are operating at their own frontier as they grapple with big questions like how to govern systems that act autonomously, to how to build resilient institutions that can adapt to an era of rapid technological change.

This report lays out emerging opportunities and challenges through the real experiences these leaders are living in 2026. It offers key reflections presented thematically and without attribution, based on what we have learned listening closely as the field evolves.

Methodology

CCF supports state, territorial, and tribal governments as they navigate the rapidly evolving capabilities of AI and other emerging technologies. We facilitate a national community of practice that includes more than 40 jurisdictions, along with working groups, convenings, and ongoing one-on-one dialogue with Chief AI officers, health & human services secretaries, labor and workforce leaders, and governor's office strategists.

This report synthesizes insights from working groups, monthly and annual convenings, in addition to dozens of conversations with state and territorial government AI leaders. All conversations were conducted under Chatham House rules, and insights and quotations are shared without attribution to any individual or state. Any named state examples are drawn from publicly available materials, linked on our **AI Knowledge Hub**.



PART ONE
Building the Foundation

The first wave of AI needs to be focused on building institutions that can adopt them responsibly, with early wins for states that have done so.



Adaptive Governance Creates Room For Innovation

Robust organizational capacity and decision-making infrastructure are the foundational governance structures on which innovation and strategic flexibility are built. States that are further along in their AI adoption have solid governance foundations that can support and scale a variety of AI-enabled system deployments, while retaining safety and privacy protocols.

One state built the machinery for intake, triage, governance, and monitoring early, and it now runs close to autopilot. That foundation is why the state can support hundreds of use cases through its Center of Excellence.

Massachusetts negotiated a hard-won enterprise AI contract that puts approved tools in the hands of staff across state agencies, without each agency running its own procurement. Other governments can buy through that same contract, extending the reach of one negotiation well beyond Massachusetts. Another state published a comprehensive policy and is rolling an agent library out to its first agencies in a matter of weeks. Each of these wins lowers the cost of the next one and starts to make exploring more complex AI use cases feasible.

As states move from isolated pilots to broader scale deployments, the next frontier of innovation depends on developing the right continuous monitoring strategies for these systems once they go live.

Over the next year and more, we're focused on assisting states shift from static, document-based governance to continuous, operational governance as early pilots become permanent infrastructure. As one of our partner states framed it, AI governance will eventually look more like cybersecurity than policy compliance. Our aim is to continue to collect and share best practices in runtime monitoring, continuous auditing, and incident response that help states shift away from their reliance on rather than design-time intake and approval.

Change Management Is a Core AI Capability

Along with governance, a definitive change management plan represents a cornerstone of a state's AI strategy. In an era of markedly low public trust in government to deliver effectively, a definitive change management strategy tackles two major challenges states currently face:

1. People increasingly deal with private services that use agentic AI well, and they bring those expectations to government.
2. An experienced benefits workforce is watching this technology arrive without knowing how it will change their daily work, or whether it will cost them their jobs.

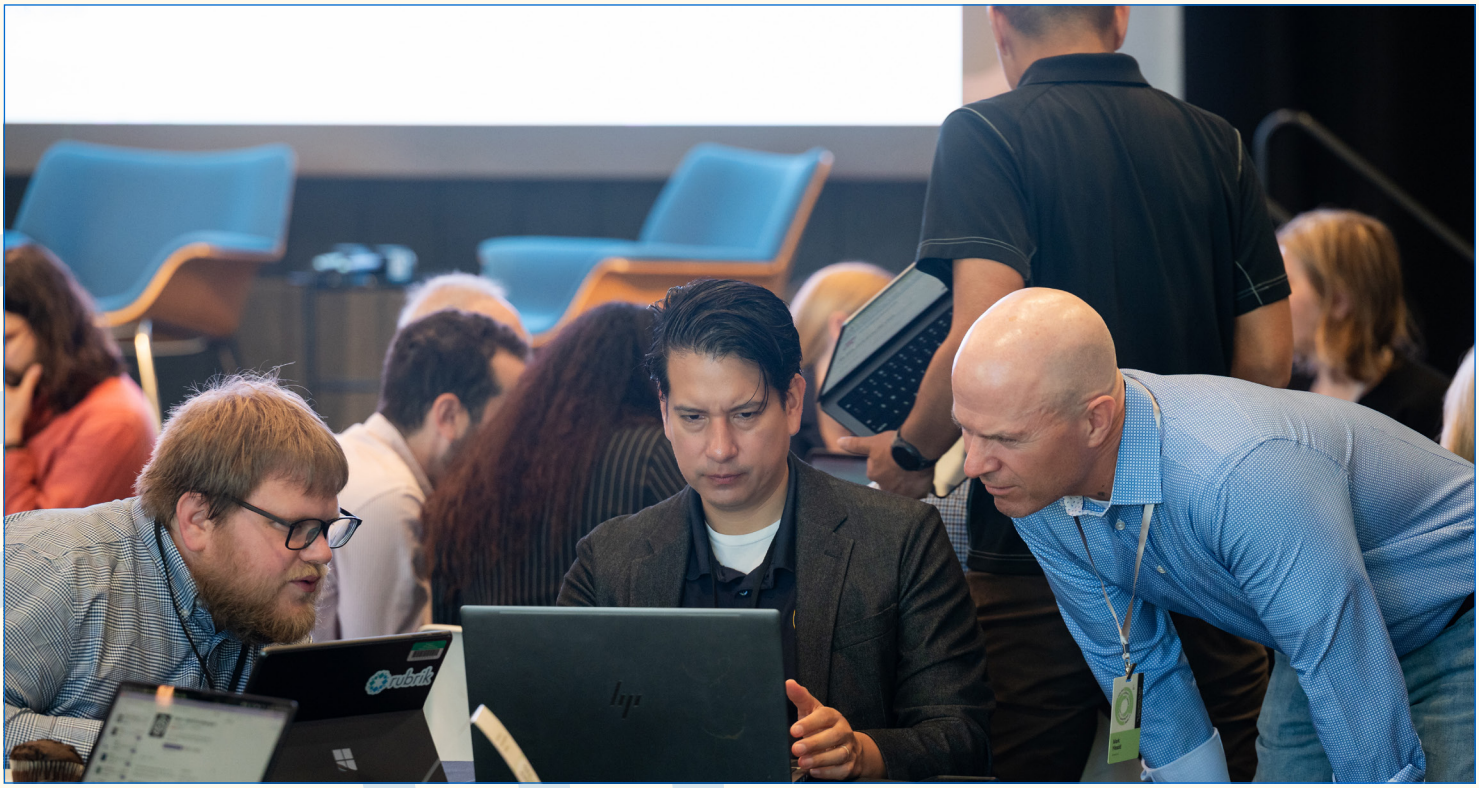
The latter audience especially is front-and-center for state leaders that we spoke to in our community of practice.

Longtime public servants have lived through waves of technology rollouts that made big promises, but failed to make life easier for staff and beneficiaries. While administrations and their priorities turn over, the core staff who absorb each rollout remain. Their skepticism is rooted not just in resistance to change, but in institutional memory.

AI is more complex than previous waves of technology because it also carries a pervasive public narrative that it will eliminate jobs. Leaders who ignore that narrative, or wave it away, find adoption stalling no matter how sound their governance frameworks are.

As the pace of technology continues to accelerate, bringing people along will become core to successful government. This is why change management isn't a nice-to-have, but a core capability that government must adopt.

Change-management approaches that treat front line skepticism as a design input and a way to document how change actually happens in state government are crucial. This muscle must grow stronger with each initiative rather than resetting with every administration. Because generative AI tools provide staff with the opportunity to experiment directly, it is more possible than ever to build innovation mindsets from the ground up, create safe spaces for experimentation, and involve the people doing the work from the very beginning, rather than as a communications step at the end.



A Rise of AI Coding Tools

The number of applications for AI in government has ballooned over the last few years, with AI coding tools acting as a genuine accelerant, democratizing who codes, building tools more broadly across government and increasing the rate of user need for iteration.

One of the clearest near-term wins is AI coding assistance, which is speeding up engineers, helping staff switch programming languages, and standing up working prototypes in days.

Maryland has gone further than most, publishing an AI governance card for AI coding assistants. This guidance applies to every executive branch agency and classifies tools by autonomy level while setting a clear standard for accountability. Under this guidance, developers must be able to explain every line of code, and whether they wrote it or AI did.

While these coding assistants anecdotally come with time-saving wins, they also carry the risk of technical debt. Apps built quickly without human engineers often lack logging and observability, creating deferred maintenance cost. At least one state has paused rollout to development staff while it works through these dynamics.

With appropriate guardrails implemented, some states have found the burden of maintenance mitigated, and kept real and sustained productivity gains across the whole lifecycle.



PART TWO
The New Frontier

As AI becomes more operational, states are encountering a class of challenges they've never seen before. Properly diagnosing these emerging challenges could mean the difference between a successful or unsuccessful AI adoption strategy.

AI Agents Are Arriving Faster Than Governance Can Keep Pace

While the public conversation about agentic AI in government tends to be a few steps ahead of where states actually are, it's up to states to both define what it means and how implementing these technologies will affect the public and their workforces. Doing that groundwork will help states build the trust they need to deliver on the AI adoption plans they develop.

This work is ongoing, as many states still haven't agreed on what "agentic" AI means, but face a shifting technological environment that demands new policy regardless. For instance, one senior AI advisor we spoke with has decided to strip the term "agentic" from their materials entirely, worried that it is standing for too many things at once. Meanwhile some mature states are starting to explore agent deployment with early positive results.

CCF's agentic AI working group continues its work to develop definition and scope for what a good agent candidate looks like; developing materials that answer questions of agentic AI that remain rooted in change management fundamentals: Who is responsible when an agent acts? How do you measure whether it's working or that there is a strong return on investment? What does agent identity mean inside a government system?

Answering these questions and sharing their governance approaches across jurisdictions will provide the clarity and buy-in for the public and state workforces to leverage agentic AI implementation initiatives, rather than feel left behind by them.





States Are Outgrowing the AI Use Case Inventory

Many leaders are united around a weary frustration with the AI "use-case inventory" — the legally or administratively mandated list of every way a government uses AI. Several states are putting energy into maintaining one, but the benefits and purposes long-term are unclear.

Early on, use cases serve a real purpose of demonstrating accountability and building awareness of what is being deployed. These inventories become less useful the more a state matures in its AI usage.

To make better use of staff time, policies around use case inventories should not dictate that they be maintained in perpetuity, but reevaluate their value after a set amount of time.

AI teams with the largest deployments say that maintaining a live, uncurated inventory is simply not sustainable at scale. They are shifting their attention from cataloging projects to embedding AI in service delivery and watching how it behaves in production. The use case inventory may be seen as a helpful tool for states early in their AI journey but might not be a valuable long-term best practice for the entire field.

Government Isn't Built For Consumption-Based Pricing

Over the last couple of years, the dominant anxiety in government AI was risk around issues like bias, privacy, and hallucination. While this is still present, it's the cost and pricing structure of these AI models that states are now concerned about.

Specifically, states are worried about the move from predictable annual licensing fees to consumption-based, per-token pricing that can spike without warning.

One state lead described running a single 300-page document through an agentic workflow. It took about 15 minutes and "burned a lot of tokens," the cost of which could seriously impact a department's operating budget. If running one document through a workflow costs that much, then how thoroughly you check anything becomes a budget decision. Officials are openly asking how to optimize their token use through better prompting and smarter model selection, treating compute as a limited, highly-valuable commodity.

The deeper worry is financial sustainability. In one state, work on a promising prototype stalled when outside support for the compute costs ended. Others described the challenge of budgeting systems built for waterfall appropriations that can't accommodate on-demand pricing at all.

Some states are responding by building cost visibility directly into their platforms. One state's internal tooling lets teams see their token usage and what the project will cost them per month. But that level of instrumentation is rare. For most states, the question of what happens when an agentic system exhausts its budget midway through serving a resident remains genuinely unanswered, making states apprehensive to scale. This is an appropriations problem showing up as a technology problem.

Solutions range, but CCF will continue its work on exploring and building on shared models developed in discussion with AI officers and universities, exploring project costs and proposing strategies that could help budgeting in a world of new pricing models.



PART THREE
Shaping What's Next

The next phase of AI readiness will require shared infrastructure and new ways of working. We give a glimpse on big questions we still need to answer and the big opportunities on the horizon.

States Are Acting Like Fifty Buyers Instead Of One Market

State government's biggest untapped advantage is its collective purchasing power.

Fragmented state-by-state purchasing agreements for AI tools that most states all utilize similarly has resulted in higher costs for everyone.

Cross-state alignment on shared procurement structures and reusable hiring vehicles is emerging. Cooperative purchasing agreements—a long-standing tool in states' toolbox—can pool negotiating leverage that one state alone may not have. It also lets other states receive the benefit of existing contract terms rather than starting from scratch, providing tools, support structures, and pricing models that are better shaped to the unique needs of government systems and workers delivering services to everyday people.

Models will vary, depending on specific state needs, geography, and other factors. But embracing this collaborative approach to AI purchasing not only lowers costs for participating states but ultimately shapes an entire market to be more receptive to their needs and makes AI work for the public.





States Are Interested in Diversifying their LLM Access to Avoid Lock-In

AI adoption hasn't changed the core questions every state wrestles with when it comes to developing key service delivery products: Do we build custom AI capability in-house? Do we buy it from vendors? A blend of the two?

What AI adoption has provided though is the opportunity for states to approach purchasing from a new perspective: abandoning the old "enterprise-model" with its multi-year, locked in contracts, and building a diversified portfolio of LLMs based on a model-agnostic platform infrastructure that is responsive to the real-time needs of government departments.

States that lock themselves into a single application (Claude, Copilot, Gemini, OpenAI) inherit that vendor's roadmap and pricing, while states that build model-agnostic platform infrastructure keep their portability. One AI Officer flagged, "A lot of people are still talking at the application level and I think that is a challenge."

Especially in the ever-changing AI market, with new features and use cases being created every year, states need to maintain the flexibility in their contracting to meet new demands and needs.

Of note to the states drawn to building their own LLM, which has a lot of pros, especially in terms of cost savings, the hard part is attracting and retaining the tech talent to sustain what they build. There is no one size fits all solution but states that work together, diversify and consider their talent potential as part of their approach will be better positioned to stay less beholden to private sector markets.

Ultimately, we want to look beyond the question of build vs. buy, to a model-agnostic platform that keeps states from being locked into a single application. At the same time, those interested in the benefits of build are wrestling with the challenge of keeping top talent, and make that a critical barrier to advancing AI adoption.

The Next Phase Of AI Readiness Requires Shared Standards

States remain uniquely poised in this political environment to be at the forefront of how public sector AI adoption is developed not just nationally, but across governments worldwide. However, a key prerequisite to any comprehensive government AI readiness plan would be the adoption of shared standards, especially in the areas of environmental impact and workforce AI literacy standards.

On AI literacy, resoundingly state AI leaders say workforce training remains one of the least mature aspects of their AI readiness. Governance frameworks and risk assessment processes exist, but structured up-skilling beyond a short training videos, mostly does not.

The deeper problem is that the field cannot agree on what strategies best support AI literacy, be it office hours, educational platforms, credentialing, or embedding development teams or something else. One state AI leader shared an emerging tiered-credentialing proposal — foundations to advanced to governance — and immediately named the risk, "all these different states will likely come up with their own licensing or credentialing, and then how does that match between states? It gets quite messy."

A handful of states are independently evaluating the same handful of training platforms and arriving at the same tiered models, with no mechanism to align them or evaluate their success. A shared, portable literacy framework with agreed definitions, common tiers, and group purchasing would spare 30 states from solving the same problem 30 different ways.

Meanwhile, states are fielding legitimate questions from staff and the public about AI's environmental footprint, and lack the data to answer them. They are being asked about AI's energy and water footprint and they cannot answer questions about trade-offs clearly.

The problem is that there is no per-token energy or water metric. The frontier labs don't publish one, so there is no standardized way to assess the footprint of a given workload. It is a live accountability question with limited data behind it and it grows more conspicuous as compute use scales.

States are interested in having AI companies publish per-query energy and water data to create more transparency. Until they do, states are exposed to a question they won't stop hearing. On the whole, a shared, honest "here is what we can and cannot measure" position would serve states better than silence.

Final Note: Driving Toward the Future

These takeaways represent a complex landscape of recommended tools, proposals, and best practices, all of which are designed to put states in the driver's seat of their own AI implementation strategies.

Our commitment at CCF is to keep holding space at the table for the leaders making the decisions, the public servants doing the work, and the residents who rely on services and supports to help their families survive and thrive. From this we can turn what the field learns into things the field can use.

We're doing this by regularly convening our 35+ state AI Advisor Community of Practice, providing a candid, secure space for high-level AI implementers to share stories, best practices, and strategies. We're also sharing findings with the broader field through our **AI Knowledge Hub**. Additionally, our AI workgroups delve deeper into a variety of topic areas covered and produce comprehensive, shared frameworks that help the entire state AI Readiness field.

We also fund the $20 million **Public Benefit Innovation Fund**, supporting the development and deployment of AI-enabled technologies that improve access to essential public benefits like Medicaid, SNAP, and child care. These funded pilots and projects help lay the foundation for state AI readiness initiatives across the country.

If you lead this work in a state, territory, or tribal government, we'd welcome you into the community of practice. If you fund or build in this space, this report is your map of where support is needed most. Please reach out to **jessica@centerforcivicfutures.org** or visit **centerforcivicfutures.org** for more information.