ASPEN POLICY ACADEMY
aspen institute

Policy Brief

Implementing the Novel COMPASS System for an Adaptive AI Policy Review Process

By Erika Cheung, Amy Koo, Jordan Loewen-Colón, and Kash Sridhar
November 2026



This policy project was developed by Aspen Policy Academy alumni. They worked with the Georgia Technology Authority to help identify gaps and possible updates for the state's enterprise artificial intelligence use policy. This policy brief provides an overview of the team's proposed solution for states like Georgia seeking an AI policy that keeps up with technological advances. The full project is available here. Please note that all authors' opinions published here are their own. This publication does not reflect the views of the Aspen Policy Academy or the Aspen Institute.

Executive Summary

To proactively ensure that artificial intelligence (AI) policies, standards, and guidelines stay relevant as technology evolves, state agencies should adopt the novel Comparative Oversight and Monitoring of Policy Across States and Sectors (COMPASS) system, which uses a gap analysis tool to surface potential updates. AI tools can be safely used to automate the most time-consuming elements of the policy process (e.g., collation of research, text transcription, and categorization) and help overburdened and underresourced departments stay relevant and effective, so long as humans remain in the loop to review the results and finalize the changes. The COMPASS system would allow state agencies to adapt to the rapid pace of AI technological advancements, avoid policy drift, and improve their ability to meet their own needs and those of residents.

Background

Many state AI policies run the risk of falling out of alignment with new laws and technological developments because of the rapid pace of change. We partnered with the Georgia Technology Authority (GTA), which asked us to help it update its enterprise policy (PS-23-001) and standards (SS-23-002) related to the responsible use of AI within state agencies. GTA's current AI policy was developed in 2023. Since then, significant advances in AI technology (such as agentic systems and token-based services), along with the adoption of new laws and regulatory requirements, have created a need to revisit and

update the framework. For example, GTA's current AI policies on human review may not be aligned with Georgia Senate Bill 444, which is set to go into effect on January 1, 2027. The law states that a health insurer can't deny a claim based on a computer alone, and a medical professional has to be involved. GTA's current AI policy only requires review and approval "by the business owner or designated personnel" (SS-23-002 §5). Georgia would, in effect, be holding insurers to a standard it doesn't hold itself to.

Georgia is not the only state agency dealing with policy drift as a result of AI innovation. Governments at every level face a shared dilemma in keeping AI policy, standards, and guidelines current. One challenge that states face is the sheer amount of work and inventory needed to examine their AI policies and their impacts on and intersections with other state laws and regulations. In many cases, the labor needed to accomplish this time-consuming work has required the creation of entire offices and dedicated support staff. California, a large state, has a Department of Technology with about 1,100 positions that can absorb the workload of policy revisions. However, many smaller states and states with fewer resources face limitations in their staff, budget, and infrastructure to generate and maintain a system that keeps pace with AI innovation. This can make regular policy review and implementation updates unsustainable. The manual nature of the work has prompted some states to consider using AI tools to help keep pace.

Though there is some reluctance among constituents around AI use, many governments are already leveraging AI tools in specific contexts to solve similar problems. State and municipal governments in Ohio, Virginia, and San Francisco have been using other AI tools to help reduce regulatory burden for a number of years. In particular, AI tools are useful for scanning existing state provisions and identifying outdated, duplicative, or conflicting regulations. When developed and used properly, they are as capable as, and sometimes more comprehensive than, humans in deciphering, restoring, and interpreting large and complex texts. These AI tools are proving that, with the right guardrails, lawmakers and state regulators can be more efficient and better serve their constituents.

Recommendations

We worked with GTA as a use case to help identify where the agency's enterprise AI policy was outdated using a novel gap analysis tool (assisted by AI). We used the same tool to recommend updates to GTA's policies (which would ultimately be reviewed by the agency's state policy team for alignment, accuracy, and compliance). We recommend other state agencies should implement the COMPASS system to help compare their AI policies with other benchmarks and to set a direction for future revisions to AI policy. This process includes implementing a novel technology assisted gap analysis and monitoring

tool and reviewing policies both on a fixed schedule and whenever key legal or technological changes occur.

The COMPASS system combines human input with automated assistance. The gap analysis tool can compare a state's current policies to best-in-class policies around the world; propose detailed updates for staff to review, including exactly where changes (additions, removals, adjustments) need to be made in current regulations; and generate sample markup language. For example, the tool may recommend updating language around AI risk and governance terms or help set safety control language for agentic AI. Policy staff can then select which updates they would like to implement. For instance, the tool proposed 36 policy updates for Georgia, which we narrowed to 12 final policy updates for the GTA staff to review.

The value of an AI-assisted review process is that it leverages the strengths of contemporary AI tools (finding, extracting, comparing, and drafting) while avoiding their weaknesses: judgment and decisionmaking. Automated tools can be carefully designed to scan authoritative sources (decided by humans), extract requirements and citations, create comparison charts of an external approach (carefully tuned to an agency's own language), and suggest changes, with every claim traceable to a human-verifiable and human-approved authoritative source. To account for the risks associated with using automated AI tools – namely, the possibility of AI hallucinations or false outputs – the COMPASS system keeps a human in the loop for ultimate verification and judgment.

The COMPASS system should be implemented in 2 phases:

Phase 1: Implement a policy review process using a technology enhanced gap analysis and monitoring tool and human review

State agency policy teams should develop a semiautomated review process for their AI-related policies, standards, and guidelines that can be integrated into their workflow and infrastructure.

The review process is a system to flag legal, regulatory, and technology changes across a specific locality, the United States, and major AI safety research organizations that may affect an agency's current policies. With the help of a gap analysis tool, policy teams can analyze existing policies using a 5-step process:

1. **Define direction:** The policy team confirms authoritative source texts in need of updates and names the priority topics (e.g., security, definitions, and scope).

2. **Research and compare:** The automated gap analysis tool analyzes curated sources, pulls out key requirements and citations, and builds a comparison chart against existing policy; the owner checks whether that evidence is credible, relevant, and complete.
3. **Validate and prioritize:** The team validates which evidence matters, ranks the gaps and risks, and rescopes.
4. **Draft and refine:** The tool collects and refines the evidence gaps into a reviewable draft package.
5. **Approve and publish:** Suggested policy updates are brought to a policy committee for final review. The team edits and approves the final language; completes governance, legal, security, and leadership review; and authorizes publication.

Once the COMPASS system is built, it becomes a real-time monitor of shifts in legislation and technological developments that may conflict with current policies, guidelines, and standards. Each run of the gap analysis tool is tracked and used to maintain traceable changes and updates through every subsequent run.

Phase 2: Create a calendar and event-based policy revision process

The COMPASS system should be run semiannually based on the policy team's state governance rhythm. The system should also have an event-based trigger that is determined by technology and legislative developments.

For example, GTA's 2023 policy and standards did not become outdated at a single point in time; they became outdated when new capabilities (e.g., generative AI, agentic systems, and token-based cloud services) arrived between review cycles. A review process tied only to a calendar date will always trail AI's advancements. Therefore, agencies should pair scheduled reviews with specific triggers for additional review, such as new laws or major technology advances, such as the following:

- Enactment or amendments of a state law addressing AI;
- A federal rule or regulation reaching state agencies with a compliance date;
- A judicial decision materially affecting state use of automated systems;
- A material change by a supplier to a system in state use (e.g., AI model, version, subprocessor, newly activated capability);
- An AI-related incident in the policy team's state or a significant one in a peer state; or
- Emergence of an AI capability that current instruments do not address (e.g., agentic systems in 2023).

An example of the COMPASS system's output is shown below. This prototype was shared with the GTA team for their review; it does not constitute actual legal or regulatory language and would ultimately need to be reviewed and adopted by the state to have effect.

AMENDMENT CARD 7 OF 12 · SS-23-002 (STANDARD) · NEW §14 — AGENTIC OPERATIONS

Set Safety Controls for Agentic AI

Default to no autonomous state action without scoped human authorization.

TODAY

Current guidance says AI should not have autonomous decision-making capability but no binding control governs agentic operation. An agency employee or resident can encounter this gap without a testable record.

AFTER THIS AMENDMENT

No AI System performs an agentic operation in a State environment unless a human principal has authorized the exact class of action, the system is technically constrained to that authority, and logging, interruption and testing requirements are satisfied.

PROPOSED TEXT (ADD)

Where: New §14 — Agentic Operations · **Change type:** New provision · **Replaces:** "New provision. Current guidance says AI should not have autonomous decision-making capability but no binding control governs agentic operation."

SAMPLE

14. Agentic Operations.

(a) Default rule. An AI System shall not perform an Agentic Operation in a State environment unless a human principal has authorized the exact class of action, the system is technically constrained to that authority, and the requirements of this section are satisfied. Producing an output for human review is not an Agentic Operation unless the system also takes an action.

(b) Before enabling an Agentic Operation, the Agency CIO and GTA OIS shall approve a written authority profile specifying: human principal; permitted tools and actions; prohibited actions; data and systems accessible; credential scope; financial or transactional limit, if any; geographic or program scope; duration; delegation rule; required confirmation points; interruption method; monitoring owner; and termination condition.

(c) An agent shall use a distinct identity or otherwise produce logs that reliably distinguish the agent's actions from the human principal's direct actions. Shared human credentials shall not be used to conceal agent activity.

Agentic Operations (Continued)

(d) The system shall not delegate authority to another agent, model, vendor service, or subprocessor beyond the approved authority profile. Cross-vendor delegation is prohibited unless expressly approved and contractually governed.

(e) A High-Risk Agentic Operation shall provide an effective means for an authorized human or GTA OIS to pause or terminate the running process. If the Agency cannot demonstrate interruption before deployment, the operation shall not be enabled.

(f) Logs shall be sufficient to reconstruct material actions and shall include the human principal, instruction or authorization, system/model version, tool calls, material retrieved content, external communications, data or configuration changes, delegated calls, time, result, and interruption or override event.

(g) Before deployment and after a Material Change, the Agency shall test direct and indirect prompt injection, unauthorized tool use, privilege escalation, cross-user disclosure, prohibited delegation, and interruption. A successful unauthorized action or inability to interrupt is blocking.

(h) Long-lived authorization shall expire no later than thirty days after issuance unless a shorter period is required by the risk assessment. Renewal requires review of logs, incidents, changes, and continued necessity.

(i) Evidence shall be retained in the security repository and statewide AI inventory. Agentic operation outside the approved profile shall be stopped immediately and treated as a Reportable AI Incident and noncompliance under section 11.

Owner: Agency CIO and GTA Office of Information Security; named human principal · **Findings:** BASE-2026-01-F03, F26 · Proposed text is complete and needs to be reviewed by a human prior to adoption.

Companion Resources

We have provided additional resources to give greater context, including details of how the COMPASS system was built and a video explaining the demonstration and walkthrough of how the system could be used for GTA.

- **COMPASS system plan:** A detailed explanation of how humans and AI engage as part of the COMPASS system.
- **Demonstration video of policy updates:** An explanatory walkthrough video for GTA and a refined output of the COMPASS system.

FAQs

What about other states?

Georgia is not the only state dealing with policy drift due to AI innovation. Governments at every level face a shared dilemma in keeping AI policy, standards, and guidelines current. Two pressures compound one another. The legal environment is moving quickly: States introduced fewer than 200 AI-related bills in 2023, but in the first quarter of 2026, lawmakers in 45 states had introduced 1,561 bills. Making things even more difficult is that federal guidance is unstable; the Office of Management and Budget's 2024 memorandum on agency AI governance was rescinded and replaced in April 2025. Meanwhile, generative AI has raised new questions about government use of AI tools, prompting states such as Georgia to supplement 2023 anchor policies with dedicated generative AI standards. The result is a need for more frequent revisions to address increasingly technical subject matter.

New York State offers the clearest example of transparent revision: Its Acceptable Use of AI policy has been updated twice since it was issued in January 2024, with each change logged in a public revision table. The most recent revision is owned by a chief AI officer, a role first filled in March 2025, to refine the state's AI guidance. California has taken a different approach, distributing its generative AI requirements across multiple administrative and IT manuals maintained by the Department of Technology, a department of about 1,100 positions whose policy staff are part of a much larger operational organization. Many states are unlikely to have comparable staff capacity, investment, and infrastructure to sustain this pace of revision on their own.

f11photo via Getty



However, states that are trying to keep pace have not relied on staff size alone; they have built specific governance structures. Washington State developed its statewide AI policy through a multiagency and local government subcommittee and documented the revisions, which turned interim generative AI guidelines into a broader policy. Maryland paired an AI subcabinet, later backed by statute, with a dedicated AI Enablement team inside its IT department. Vermont, one of the smallest states, created a statutory Division of Artificial Intelligence in 2022, anchored by a single authorized position and an advisory council. These cases suggest that the binding constraint is less raw headcount than whether a state has designated an accountable owner, a routine revision process, and a mechanism for cross-agency input. Many states have not yet made those investments, and state chief information officers continue to report funding and governance challenges even as AI has become their top priority.

"State agencies should implement the COMPASS system to help compare their AI policies with other benchmarks and to set a direction for future revisions to AI policy."

Erika Cheung, Amy Koo, Jordan Loewen-Colón, and Kash Sridhar

How can AI policy gaps prevent AI enablement in state agencies?

Beyond addressing specific risks, clear AI policies can also help state agencies identify opportunities to use AI to improve services. For example, Georgia Pathways to Coverage is a state Medicaid program that provides healthcare insurance to low-income residents who do not qualify for traditional Medicaid. To receive this benefit, residents must log 80 hours of work, school, or volunteering each month. Between 2020 and 2025, the Georgia Pathways to Coverage program spent $80.3 million on the program; $54.2 million was administrative costs. Additionally, as of August 31, 2025, more than 9,000 people were enrolled in the program out of an eligible population of about 250,000. A report by the US Government Accountability Office identified enrollment platform glitches, misunderstandings, and administrative complexity as contributing challenges.

Why not just ask generative AI models to update the policies themselves?

The speed and scale of AI development demonstrate that we need a new approach to policy revisions that leverages the unique attributes of the technology. However, many people are wary of using generative AI in government work, and rightly so. For example, in 2025, Australia paid a contractor AU$440,000 for a report with multiple errors, which were initially attributed to human error. The contractor only later admitted to using AI when more fake quotes and citations were discovered. The event quickly triggered updates to Australia's procurement policies, as vendors were not previously required to disclose their use of AI.

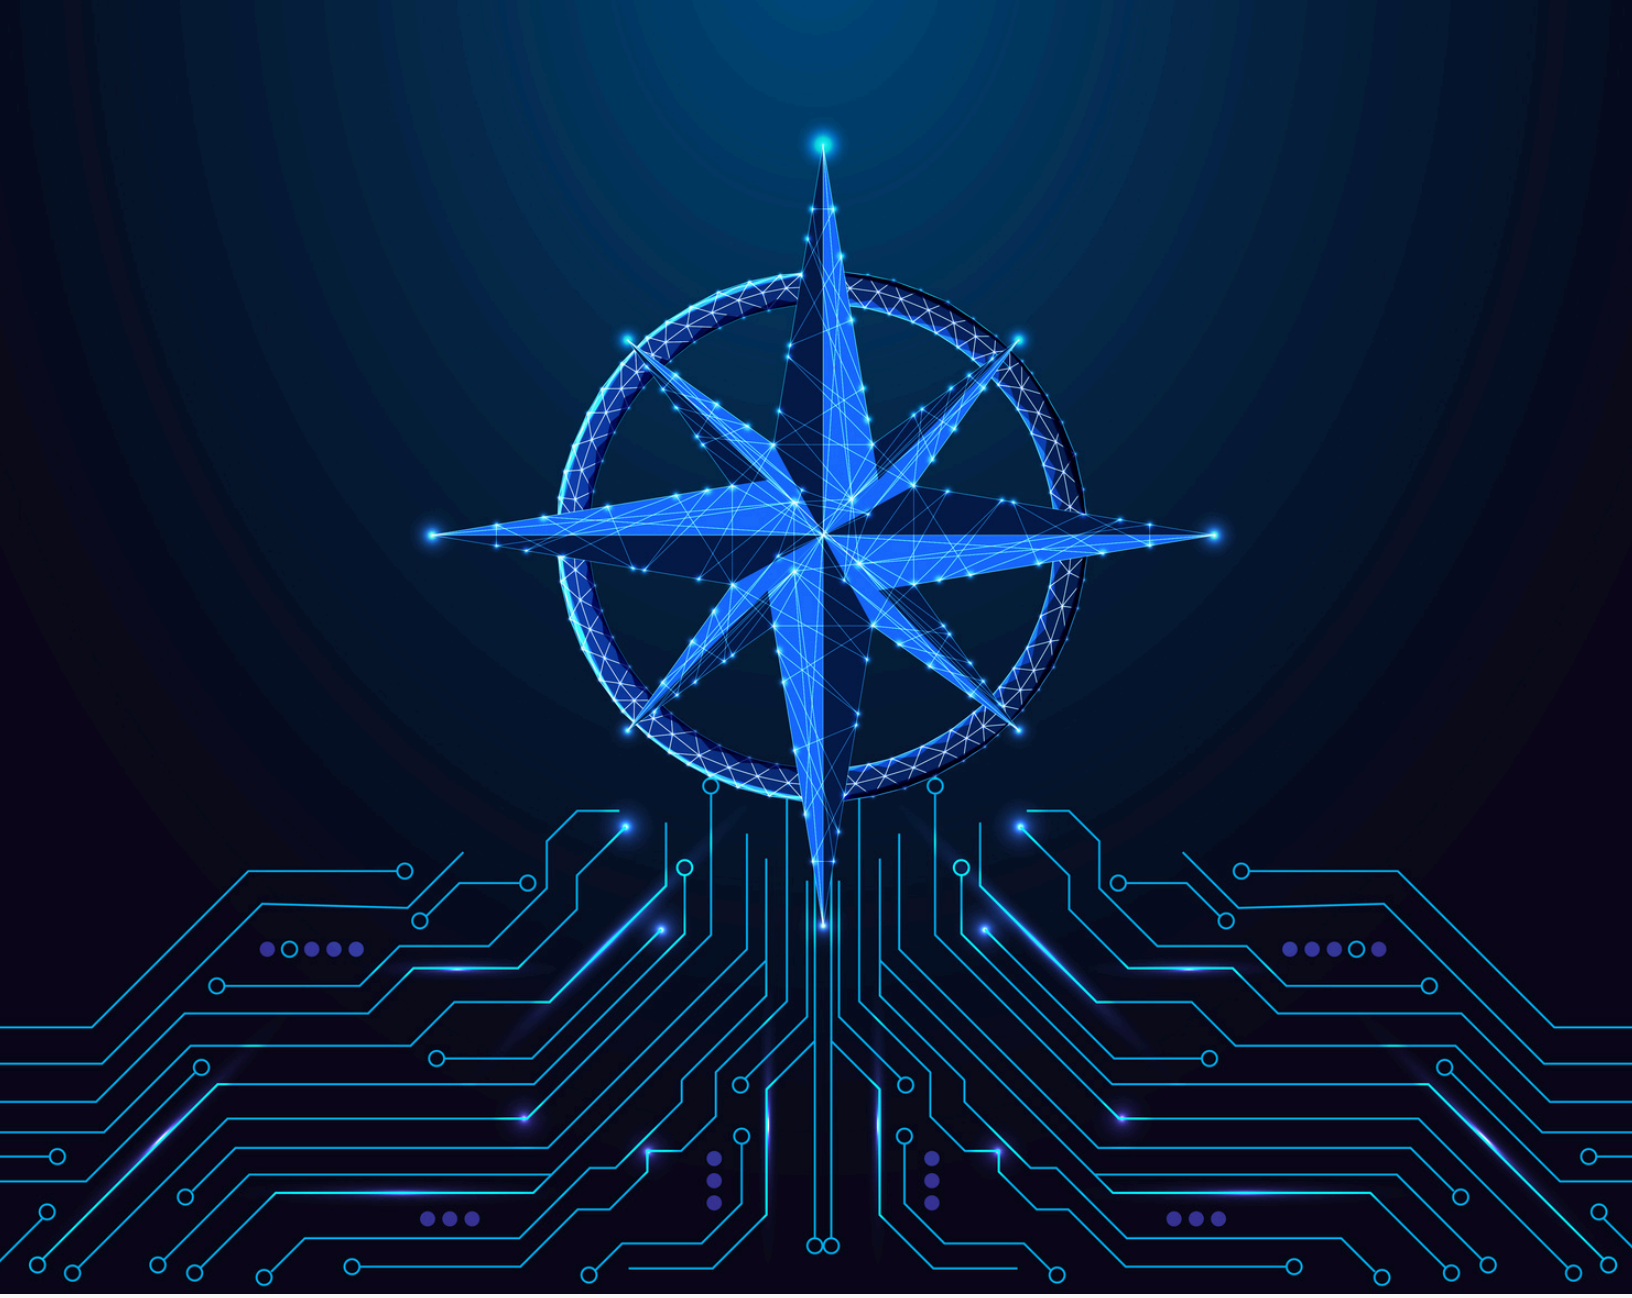
While tech failures such as fabricating quotes and citations are real, the problem is not that these tools cannot be used in policy work. It is that they cannot be used without human verification in policy work. This is why the human designed guardrails of the COMPASS system are the most important part of the automation process. Specifically, every citation, comparison, and draft output has a human reviewer and a traceable source. The AI does the machine work, not the human work.

Can't we just crowdsource through multistate coalitions?

Because innovation is outstripping regulation, there is a growing movement for governments at the local, state, and national level to connect and collaborate on AI policy development. The GovAI Coalition includes more than 900 federal, state, and local government agency members who work together to mutually address AI concerns and research in the public sector. Through working groups, they provide shared resources, standards, and advocacy.

Georgia already has 2 members participating in the coalition: the Atlanta Independent School System and the Atlanta Regional Commission. However, individual government agencies still need to tailor crowdsourced findings to their local needs and specific statutory language.

If you'd like to learn more, see the full project, including a Gap Analysis Implementation Plan and Demo Video, at aspenpolicyacademy.org/project/state-ai-policy-monitoring-2026.



inkoly via Getty Images

About the Aspen Policy Academy

The Aspen Institute's Policy Academy helps community leaders and experts across the political spectrum elevate their voices, influence key decisions, and strengthen democracy from the ground up. Our innovative training programs and resources equip people across sectors – from tech to the environment, science to civic engagement – with the skills to shape critical policy efforts. Learn more at aspenpolicyacademy.org.