ASPEN POLICY ACADEMY
aspen institute

Implementation Guide

COMPASS Implementation Guide

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

November 2026

This Implementation Guide was developed by Aspen Policy Academy alumni. They worked with the Georgia Technology Authority (GTA) to help identify gaps and possible updates for the state's enterprise artificial intelligence (AI) use policy. This guide outlines the 5-step Comparative Oversight and Monitoring of Policy Across States and Sectors (COMPASS) system, which uses an automated gap analysis tool to compare the AI policies of agencies like the GTA against new laws, standards, and technology and produce a set of proposed policy updates for human review.

The full project, including a policy brief explaining the alumni's core recommendations, 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. The graphic on page 2 was created with the help of AI through ChatGPT. Generative AI is experimental and information quality may vary. The content has been reviewed, edited, and finalized by the authors.

About This Guide

Who this is for

- AI policy teams.

Time required

Depends on expertise; needs some training and technical understanding. An estimation would be:

- Approximately 1 day for setup; and
- About 20 minutes per run of the gap analysis tool for experienced policy teams, with variable additional time needed to verify and approve updates based on a team's needs.

Tools needed

- Large language model of choice (e.g., Claude, ChatGPT, Gemini);
- Internal policies in scope of revision; and
- Preferred external comparison policies.

How the Process Works

This guide outlines the 5-step process of the COMPASS system piloted for a partnership project with the Georgia Technology Authority (GTA). In each step, a person makes the decisions, and the AI tool does the time-consuming work of finding, extracting, comparing, and drafting the information. Each step produces 1 output that becomes the input to the next step.

The COMPASS System Process



Chart generated with AI

Ground rules for every step

- The AI only works from sources that a person has approved. It should not add outside information.
- Every claim must trace back to a source that a person can open and check.
- Nothing moves to the next step without the sign-off listed in that step.
- Human judgment and decisionmaking are crucial throughout the process.

What the Final COMPASS System Output Looks Like

The COMPASS system ultimately produces a set of policy cards based on the gap analysis tool – 1 for each proposed policy update, that the policy team can review before finalizing the updates.

In the GTA pilot partnership, the process produced 36 proposed updates, 12 of which were selected for staff review in the interactive demo. The output varies depending on the preferences of the agency and its current revision process.

Example

GTA needed to broaden its definition and scope of what constitutes artificial intelligence in order to define the system, risk, review, incident, change, and agent terms. The policy card shows the policy team the changes that could be made.

Each card answers 4 questions:

1. **Where:** The exact section of the current policy that would change.
2. **What:** Whether it is an addition, removal, or adjustment.
3. **Proposed language:** Sample markup showing the suggested wording.
4. **Why:** The law, standard, or research behind the change, with a citation.

AMENDMENT CARD 2 OF 12 · SS-23-002 (STANDARD) · SCOPE & DEFINITIONS

Define AI Risk and Governance Terms

Define the system, risk, review, incident, change, and agent terms.

TODAY The Standard defines AI, GenAI, deep learning, and machine learning, but not the terms needed by the accepted findings. An agency employee or resident can encounter this gap without a testable record.

AFTER THIS AMENDMENT The accountable owner follows AMD-2026-04; the resulting record shows whether the control was satisfied. Definitions are adopted before or with the amendments that depend on them.

PROPOSED TEXT (ADD) **Where:** Terms and Definitions · **Change type:** Add · **Removes:** “The Standard defines AI, GenAI, deep learning, and machine learning, but not the terms needed by the accepted findings.”

SAMPLE

Add the following definitions:

Agentic Operation: operation in which an AI System can select or invoke tools, call software or services, modify data or configurations, communicate externally, initiate transactions, or take another action in a State environment beyond producing an output for human review.

AI Incident: an event in which an AI System causes or materially contributes to unauthorized disclosure, security compromise, unlawful or prohibited action, material service disruption, material inaccuracy in a Consequential Determination, discrimination or material disparity, loss of required human control, or another event that triggers an applicable legal, security, privacy, or continuity process.

AI Use Case: the documented combination of an AI System, purpose, users, affected persons, data, operating context, decision or action supported, and deployment pattern.

Automated Decision System: an Automated System used to make, recommend, or materially influence a Consequential Determination.

Automated System: an information system, software, model, or computational process, including one using Artificial Intelligence, that materially assists or replaces human decision-making, judgment, communication, recommendation, ranking, classification, allocation, or action.

Consequential Determination: a decision or action that determines, denies, reduces, conditions, ranks, or materially affects a person's legal right, eligibility for or access to a public benefit or service, employment or promotion opportunity, license, contract, housing, education, health coverage, liberty, public-safety treatment, or other comparable State-provided opportunity.

High-Risk Use Case: an AI Use Case involving a Consequential Determination; law enforcement, elections administration, public safety, cybersecurity response with authority to alter production systems, or access to a Restricted Holding; or another use the State CIO determines could materially affect legal rights, benefits, safety, security, or public trust.

Human Review: a review performed by a qualified human who receives the information required by section 5, has authority to reject or change the AI output or recommendation, and records the disposition when required.

Low-Risk Use Case: an AI Use Case that is neither High-Risk nor Moderate-Risk.

Material Change: a change to the model, model version, data source, retrieval source, subprocessor, system prompt, decision logic, permissions, tools, deployment pattern, or intended use that could materially change risk, performance, data handling, or the validity of a prior approval or evaluation.

Moderate-Risk Use Case: a non-High-Risk AI Use Case that is public-facing, processes non-public State data, materially informs agency operations, or can create material financial, security, records, privacy, or reputational impact.

Restricted Holding: data or records subject to a statutory, regulatory, contractual, court-ordered, confidentiality, security, or access restriction that limits who may access, use, process, disclose, or redisclose the information.

Owner: GTA Office of AI maintains definitions; Agency CIO applies them · Findings: BASE-2026-01-F03, F06, F09, F15; F2-02, F3-02 · Proposed text is complete and needs to be reviewed by a human prior to adoption.

Handoff and Maintenance

What gets updated after each run

The policy team updates 2 anchor documents that are used to track and maintain traceability and prevent a well-maintained policy from drifting when the person maintaining it changes jobs.

Who owns it going forward

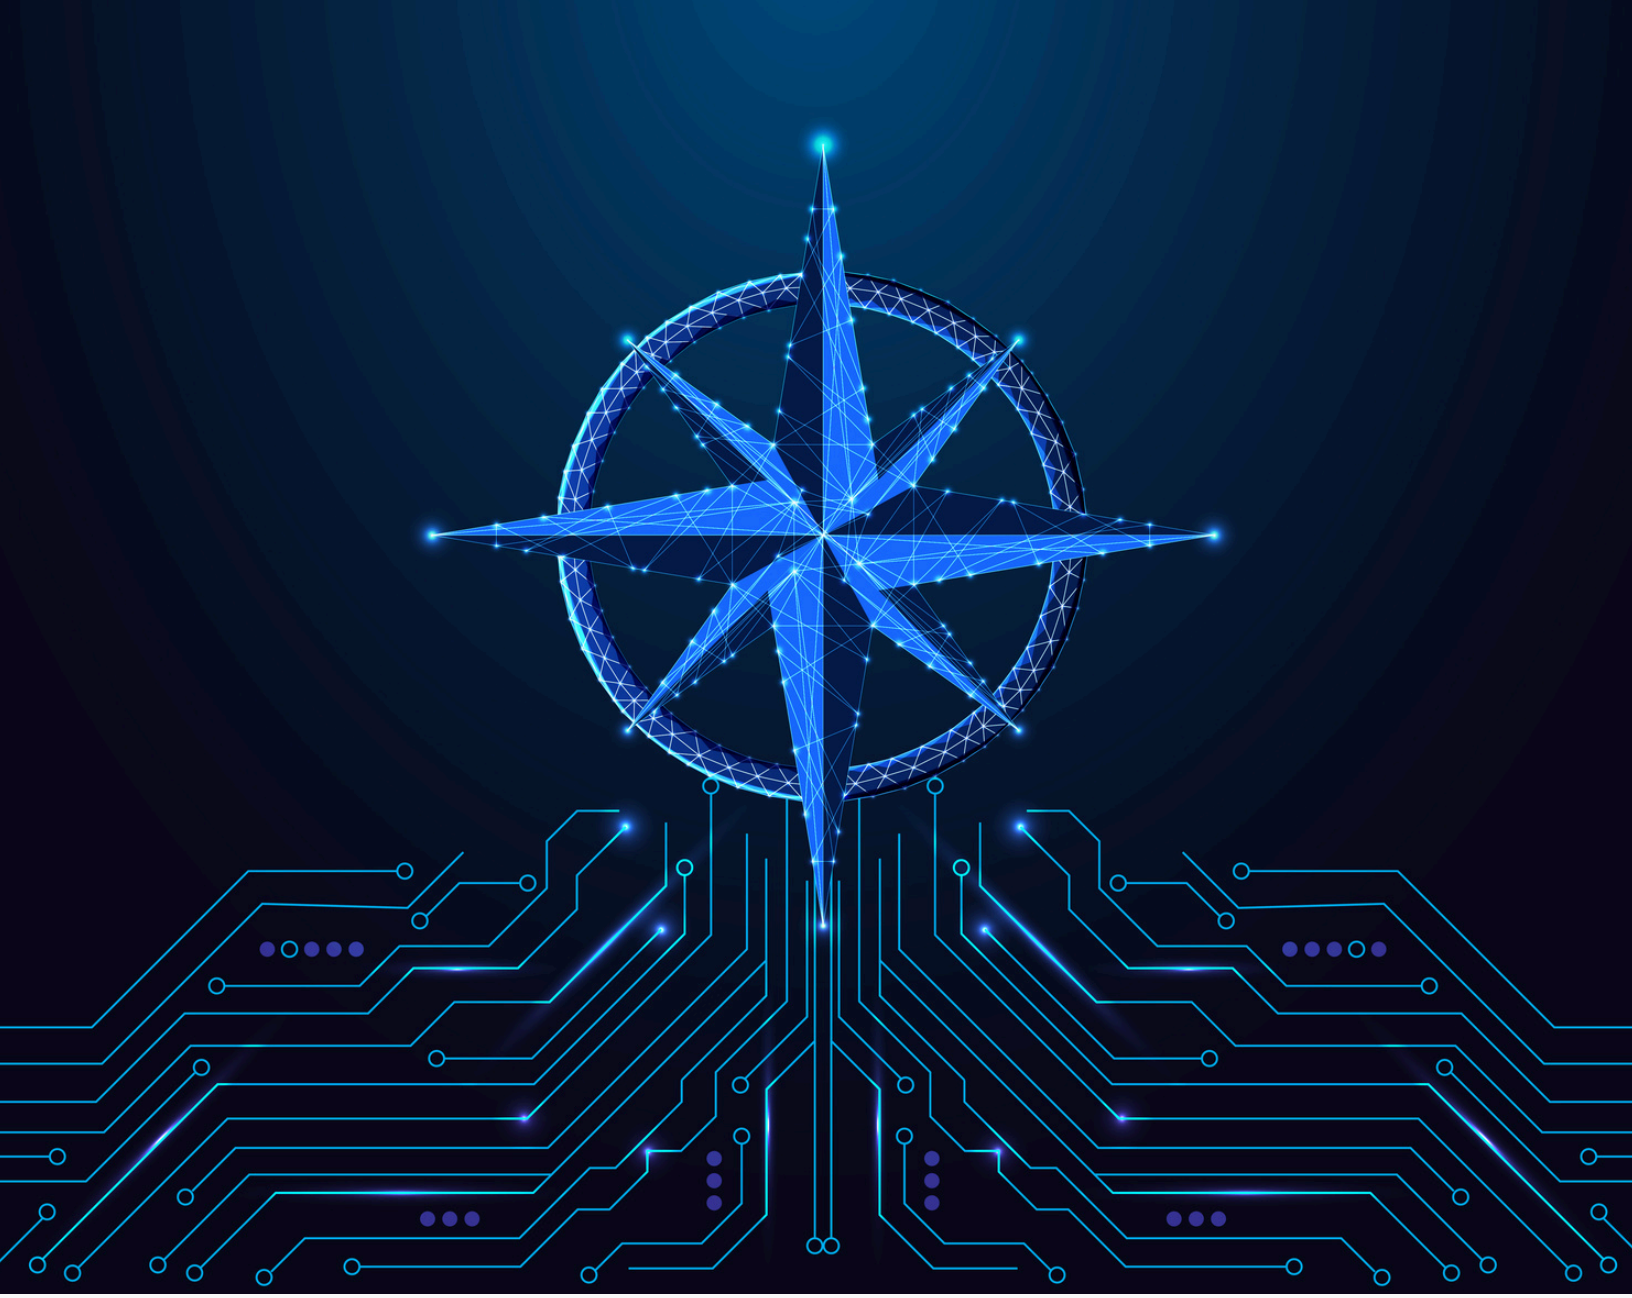
This is an iterative process, and it is important to ensure that someone continuously owns the process. To keep a well-maintained policy, more than one person should be designated as a policy owner to keep the policy from drifting when the person maintaining it changes jobs.

When to run it again

On a fixed schedule (e.g., twice a year) and whenever one of the following events occurs:

- A state law addressing AI is enacted or amended;
- A federal rule reaches state agencies with a compliance date;
- A court decision materially affects state use of automated systems;
- A vendor makes a material change to an AI system in state use (e.g., new model, version, subprocessor, capability);
- An AI-related incident occurs in the state, or a significant one in a peer state; or
- A new AI capability emerges that current policy doesn't address.

If you'd like to learn more, see the full project, including a policy brief explaining the alumni's core recommendations, 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.