ASPEN POLICY ACADEMY
aspen institute

Policy Brief

Establishing an Integrated Evaluation Process for Artificial Intelligence-Enabled Technologies

By Tiffani Martin
November 2026



This policy project was developed by an Aspen Policy Academy Fellow while participating in the 2026 Science and Technology Policy Fellowship. This policy brief provides an overview of the Fellow's proposed solution. The full project is available here. Please note that the author's opinions published here are their own. This project was created with the help of AI. Generative AI is experimental and information quality may vary. The content has been reviewed, edited, and finalized by the author. This publication does not reflect the views of the Aspen Policy Academy or the Aspen Institute. This project was developed between June 26 and August 31, 2026, with initial submission on August 31, 2026. The content reflects information available during this period.

Executive Summary

The Texas Department of Information Resources (DIR) should pilot a unified statewide process centered on the Coordinated Artificial Intelligence Requirements (CAIR) Playbook, which would require Texas public institutions to evaluate accessibility, cybersecurity, and artificial intelligence (AI) governance together before purchasing, deploying, renewing, or significantly changing AI-enabled technologies. Today, Texas public institutions review accessibility, cybersecurity, and AI governance separately, with different officials, evidence requirements, and timelines. Because no one is responsible for reconciling these reviews, an AI tool can pass each one and still reach the public with unresolved gaps; accessibility problems are most likely to be overlooked. The CAIR Playbook – a new step by step guide that shows organizations how to review accessibility, cybersecurity, and AI governance together before they approve an AI system – would address this problem by clarifying who is responsible for resolving conflicts between reviews, standardizing the evidence that each review requires, and requiring one shared decision before an AI system reaches the public.

Background

AI adoption in the Texas government is accelerating: More than a third of Texas state agencies were using some form of AI according to a 2022 DIR report, and by 2024, DIR's chief information officer reported to lawmakers that more than 100 of the state's 145 agencies were using AI in some form.

Ensuring that AI tools are accessible – particularly those used by governments – is essential to enabling all populations can use these tools effectively. If a government AI tool is not accessible, certain populations (such as those with disabilities) may be shut out of participation entirely because they cannot choose another provider for public services like benefits, licensing, or education. The number of Texans who depend on accessible AI systems is not a small population: More than 5.7 million adults in Texas – roughly 1 in 4 residents – live with a disability.

This underscores the importance of comprehensive and cohesive evaluation of AI systems used by governments for public use. Evaluating AI systems is not new, but there is currently no unified process for doing so in Texas. State agencies and public institutions in the state review accessibility, cybersecurity, and AI governance under separate rules, with different officials, evidence requirements, timelines, and approval authorities. As of August 2026, the cybersecurity team that previously functioned within the DIR in Austin has been transitioned to Texas Cyber Command in San Antonio, illustrating both the physical and figurative separation of these domains when reviewing AI tools used by the Texas government.

Texas has recently expanded its AI governance structure. In early 2026, the DIR board adopted Chapter 219 of the Texas Administrative Code (TAC), which established an AI code of ethics and minimum standards for covered government entities. Chapter 219 requires a written risk assessment before a covered entity develops, procures, deploys, or uses a heightened scrutiny AI system. A new assessment is also required after any "material change" to the system, though Chapter 219 does not define this term. A separate section of the rule covering state agency impact assessments names three kinds of material changes: changes to the system itself, to the data it uses, or to the purposes it is used for. The assessment covers security, privacy, fairness, accuracy, transparency, human oversight, accountability, and monitoring, but it does not expressly require accessibility testing, an Accessibility Conformance Report, participation by the digital accessibility officer, a formal cybersecurity risk decision, or Texas Risk and Authorization Management Program evidence.

There is currently no statewide process that consistently brings accessibility, cybersecurity, and AI governance findings together to inform decisions about a specific technology. Because Chapter 219's risk assessment does not explicitly require accessibility testing, accessibility issues may go unidentified until later stages of evaluation, when quality assurance testing assesses whether the technology meets the needs of end users with disabilities. Texas does not require these different officials to compare their decisions before final approval, and it does not say who resolves conflicts or takes responsibility for the risks that remain. When accessibility, cybersecurity, and AI governance are evaluated through separate requirements, the same technology can pass each review individually but still fail in functionality and effective use.

Texas Tech University offers an interesting model to help create a coherent process: Its policies place procurement, AI, cybersecurity, accessibility, and data sharing under one IT structure, with final authority assigned to its chief information officer (CIO). This approach has two main benefits. First, it creates a clear line of accountability from each review to the final approval, so conflicting findings are resolved rather than ignored. Second, it catches problems early. A secure system can still be inaccessible, and an accessible one can still carry security risks. Reviewing findings together before approval lets organizations fix these gaps before a technology becomes part of a contract or a public service that residents, students, and employees rely on.

Recommendation

The Texas DIR should use the Coordinated Artificial Intelligence Requirements (CAIR) Playbook to pilot an integrated evaluation process that identifies gaps and develops remediation strategies before an AI-enabled technology is purchased, deployed, renewed, or significantly changed. The pilot would test the process against real technology decisions rather than hypothetical ones, reveal shortcomings under actual workflows and timelines, and produce evidence that DIR could weigh before committing the state to broader adoption.

The officials responsible for accessibility, cybersecurity, AI governance, procurement, and the program using the technology (the digital accessibility officer, information security officer, AI risk officer, purchasing lead, and program manager) would bring their existing findings into one decision process. This would enable teams to address conflicts, assign unresolved risks, set approval conditions, and determine whether the technology should proceed. This would address the problem described earlier: Today, no one is required to compare the reviews, resolve conflicts between them, or own the risks that remain.

The integrated process would assign each of those tasks so that gaps, including accessibility gaps, could be caught before approval rather than after a tool reaches the public.

Existing reviews would continue under their current requirements; the integrated process would connect the findings at the point of approval. The new evaluation process would also name the final decisionmaker. The deploying organization, through its executive head or CIO, would retain authority to approve the technology and accept any remaining risks. DIR would define and support the evaluation process and encourage its adoption, but it would not accept risk on behalf of other entities.

DIR is best positioned to lead this effort because it already sets statewide technology policy, administers accessibility and AI rules, supports procurement, and coordinates technology planning and reporting. Additionally, because DIR already has personnel conducting reviews of AI tools, implementing this process would require minimal new costs and staffing because it would build on the existing review structure instead of creating a new one. DIR could set one common process while letting each organization fit it into its own approval structure, whether that is a state agency working within DIR's centralized systems or a university with its own purchasing authority.

Implementation

DIR should take the following steps to adopt a unified, statewide evaluation process:

1. **Charter a 12-month initiative to build a coordinated evaluation process.**
    - The executive director of DIR should charter a 12-month initiative led by the AI and Innovation Division. The AI and Innovation Division should establish a working group to help generate and implement this process and ensure that it includes staff from accessibility, cybersecurity, legal, procurement, and AI governance teams.
    - The participants should include the Texas Higher Education Coordinating Board (THECB) as the state agency participant, a university or other institution of higher education regulated by THECB as the public institution, and a local government that already has an AI governance framework in place.
    - THECB is well suited to serve as the state agency participant because these risks converge for a large student population: The National Center for Education Statistics reports that 21% of undergraduates and 11% of graduate students have a disability, and these students depend on accessible digital systems – course portals, registration platforms, testing software – to complete their education.

2. **Use the CAIR Playbook to run a 90-day pilot process with an actual or planned AI-enabled technology.**
    - The pilot should center on the CAIR Playbook, which contains 3 components:
        - **CAIR Implementation Charter.** This document puts every foundational decision in writing at the start of the process – a named executive sponsor, a named process owner, a signed approval chart, the checkpoints that the process controls, the enforcement rules, and the resources and capacity allocated to the review process.
        - **CAIR Readiness Assessment and Road Map.**
            - **Readiness Assessment.** This structured self assessment helps the organization examine its own evaluation processes; identify where they already align; and surface internal gaps to address before coordinating its accessibility, cybersecurity, and AI governance reviews.
            - **CAIR 90-Day Road Map.** This document is a practical companion showing the organization how to put the Readiness Assessment to work. It walks through each step of a coordinated evaluation: the officials and offices involved, the actions each takes, where each step belongs in the workflow, the timing for completing it, and the purpose each serves in reaching a sound final decision. The Road Map also provides templates, guidance, and action steps outlining when an evaluation of an AI system is triggered, who takes part, what evidence each area provides, how conflicts are resolved, who owns remaining risks, and who gives final approval.
        - **Measurement and Evaluation.** DIR should review the results of the pilot process and revise the playbook based on those results and participant feedback. Upon completion of the evaluation process, DIR could decide whether broader voluntary adoption is appropriate for other organizations throughout Texas.

If you'd like to learn more, see the full project, including a CAIR playbook, sample implementation charter, and CAIR readiness assessment, at aspenpolicyacademy.org/project/texas-ai-evaluation-2026.



Luke Jones via Unsplash

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.