



NASCIO State IT Recognition Award Nomination

State: Maryland

Agency: Department of Human Services (DHS)

Category: Artificial Intelligence

Project Title: AI-Powered Policy Enablement Package (APEP)

Project Dates: December 2025 (Pilot) - Spring 2026 (Statewide Launch)

I. EXECUTIVE SUMMARY

The Maryland Department of Human Services (DHS) serves as the state's primary social safety net, providing life-sustaining benefits to over one million residents. In this high-stakes environment, a caseworker's readiness is measured by two equally critical metrics: technical accuracy and emotional intelligence. A single error in calculation can delay a family's access to food or other critical resources, while a cold or bureaucratic interaction can erode public trust in government during a moment of crisis.

To eliminate the "readiness gap" between legislative change and frontline execution, DHS launched the **AI-Powered Policy Enablement Package (APEP)**. Moving beyond static documents, APEP provides a multi-modal ecosystem including synthesized audio, custom-grounded video, real-time guidance assistants, and a groundbreaking Adaptive Policy Simulator. Following successful pilots, DHS executed a statewide launch in Spring 2026. Transitioning from a centralized IT initiative to a democratized business capability, APEP now enables non-technical policy experts across the agency to build their own high-fidelity training suites in under an hour.

II. IDEA

The Business Challenge: The Crisis of Complexity

Maryland caseworkers, like in so many other states, manage staggering caseloads while navigating a constant stream of federal and state policy shifts. Historically, the agency relied on a "Policy-Out" model: distributing long, complex, legalistic memos such as policy updates, action transmittals, standard operating procedures and the like. This created a dangerous lag in "Skills-In" proficiency. Traditional training, focused on "what the law says", often fails to prepare staff for the "how it feels" reality of the frontline, where they must apply complex rules to residents in distress.

The Innovation: The APEP Ecosystem

APEP converts every new policy release into a four-part toolset:

1. Synthesized Audio & Video: Engaging, high-fidelity media for flexible, multi-modal learning.

2. Visual Knowledge Artifacts: AI-generated infographics and process flows.
3. Real-Time Guidance Assistant: A grounded "Policy Bot" providing instant citations without needing to reach out to central office colleagues.
4. The Adaptive Policy Simulator: An immersive training "sandbox" for high-stakes practice.

III. MULTI-MODAL MASTERY: AUDIO AND GROUNDED VIDEO

By incorporating Synthesized Audio and Video, APEP acknowledges that "reading a memo" is the least effective way to ensure high-stakes compliance for a modern, mobile workforce.

Synthesized Audio Overviews (The "Policy Podcast")

Using advanced Large Language Models (LLMs), DHS generates 5-to-10-minute conversational "Audio Overviews." These feature sophisticated, AI-synthesized discussions between two personas who break down the "why" behind a policy change. Caseworkers can listen to these during their commute, transforming "dead time" into active professional development.

Custom-Grounded Policy Video

Every APEP video is strictly grounded in the specific policy language and custom-tailored to the caseworker's daily reality.

- Operational Grounding: Instead of generic summaries, these videos explain exactly how a worker should implement the policy within their specific daily workflow.
- Targeted Implementation: These artifacts bridge the gap between "the rule on the page" and "the action at the desk," providing a visual guide for how the policy shifts affect verification requirements and final adjudication.

IV. THE GAMECHANGER: THE ADAPTIVE POLICY SIMULATOR

The APEP Simulator is a "flight simulator" for human interaction. It allows staff to practice both the "letter of the law" and the "spirit of service" in a zero-risk environment where errors have no real-world effect on a citizen's benefits.

Multi-Modal Interaction & Diverse Personas

Caseworkers interact via typing or live audio with "Virtual Applicants" exhibiting diverse emotional states: anxious, grateful, aggressive, frustrated, disorganized, or confused. By practicing against these varying personalities, caseworkers build "emotional muscle memory," ensuring they remain professional even when the resident across the desk is in crisis.

The Technical Grade: Compliance Feedback

Upon completion, the system provides an instant Policy Knowledge Grade. If a rule was misapplied, the AI points the user to the *exact sentence* in the official policy document that was missed and offers a targeted "re-do" to reinforce the correct application.

The Emotional Intelligence Grade: Customer Service Coaching

The simulator provides a Customer Service & Empathy Grade. The AI analyzes tone and suggests empathy-led alternatives to bureaucratic phrases, ensuring Maryland residents feel respected and heard during their transaction.

V. IMPLEMENTATION

Democratizing Innovation: From Pilot to the Lines of Business

The APEP journey followed an accelerated, evidence-based trajectory designed to move from a technical "proof of concept" to a permanent business capability in under six months.

In late 2025, the Alpha Pilot targeted a small group of participants within the Family Investment Administration. The focus was on high-stakes, high-error policy areas such as SNAP Utility Allowances and ABAWD (Able-Bodied Adults Without Dependents) rules.

After receiving feedback and making adjustments, including a library of sample questions, a Confirmatory Pilot was run in early 2026, expanding testing to Baltimore City. During this phase, we implemented "out-of-bounds" testing, ensuring that if a caseworker asked a question not covered by the grounded policy, the AI would safely state, "This information is not available," rather than speculating. The feedback was robust and positive.

A defining characteristic of the Spring 2026 launch is the shift from a technical project to a

permanent business capability.

- Decentralized Production: While the technology office built the initial pilots, the capability has been shifted to the "lines of business" (e.g., the Family Investment Administration).
- Non-Technical Enablement: Because the system uses intuitive templates and natural-language prompting, employees without any technical background are typically able to put their first complete APEP together in less than an hour.
- Sustainability: As soon as a policy is finalized, the subject matter experts who wrote it can deploy the training suite immediately, removing IT as a bottleneck.

Ethical Guardrails: Responsible AI

The project is governed by the Maryland DHS Responsible AI Policy:

- Human-in-the-Loop: AI synthesizes content, but a human expert always is presented with the official language of the policy to make their own judgement. Only a caseworker ever makes a decision on granting or denying services.
- Augmentation, Not Replacement: The tool is designed to make human caseworkers more accurate and empathetic; it does not automate or recommend final eligibility decisions.

Enterprise Alignment

APEP is not a "siloe" IT app; it is a core component of the Maryland DHS Payment Accuracy and Service Excellence strategy. By integrating AI-driven training directly into the policy release workflow, DHS has created a unified enterprise layer that ensures every one of our 24 local departments operates from a single, high-fidelity "Source of Truth."

Project Management Approach

The project utilized an Agile approach. While the technical build was iterative and rapid, content validation followed a strict Human-in-the-Loop protocol. This ensured that while the AI moved fast, the policy remained legally and ethically sound.

Key Groups and Stakeholder Buy-In

The project avoided top-down mandates. Instead, we used a Show, Don't Tell strategy. By providing caseworkers with a 5-minute AI-generated podcast summarizing a thick policy memo, the value proposition became self-evident. Marketing was internal and practical, with demonstrations and video being the main mechanisms.

The success of APEP relied on a coalition of leadership, technical experts, and frontline advocates:

Executive Leadership: Provided the vision and mechanism for cultural adoption.

Office of Technology for Human Services (OTHS): Existing staff served as the initial Architects, building the prompt engineering frameworks and securing the enterprise environment. By leveraging existing enterprise cloud and productivity suite licenses, we avoided the traditional 12-to-18-month procurement cycle and incurred \$0 in additional costs.

Policy Subject Matter Experts (SMEs): The true owners of the project. Their role shifted from "Authors" to "Reviewers," approving AI-generated scenarios to ensure technical accuracy.

Frontline Caseworkers: Provided the raw feedback that shaped the offerings.

Technical Architecture

The architecture is built on a Grounded Retrieval Augmented Generative AI (RAG) on a Large Language Model (LLM) framework. Unlike public AI, our architecture is fenced into our private departmental cloud tenant with strong privacy protections in place.

Assessment and Success Metrics

Successful implementation is defined by reducing the readiness gap.

We measure success of the tools themselves through Technical Accuracy (targeting >80% correct citations), Time-to-Deployment (targeting <60 minutes per package), and User Acceptance (targeting >75% satisfaction).

VI. IMPACT

Measurable Excellence

- Technical Accuracy: Verified 88.9% accuracy in policy guidance during initial testing.
- Operational Velocity: Reduced training material development time from weeks to less than 30 minutes for experienced policy experts.
- Fiscal Stewardship: Achieved 100% savings on new software procurement by leveraging existing enterprise licenses.
- Staff Readiness: The pilot achieved a 77.8% score for usefulness.