

RESPONSIBLE AI FOR SMALL TEXAS CITIES

By **Katryna Peart**, Municipal AI Governance Researcher and Consultant

Small and mid-size Texas cities are being approached by vendors offering AI chatbots for 311, benefits navigation, and resident services. Councils want visible innovation, and residents expect fast answers. But many municipalities operate with one- or two-person IT teams and limited budgets. In that environment, adopting AI without strong oversight creates operational, legal, and public-trust risks.

The Texas Responsible Artificial Intelligence Governance Act (TRAIGA) took effect on January 1, 2026. The law requires governmental entities—including cities—to provide clear disclosure when residents interact with an AI system and establishes statewide expectations for responsible governance and training. For many smaller cities, this marks the first formal requirement for AI oversight.

Texas has more than 1,200 municipalities. TRAIGA applies to all of them.

Why Evaluation Matters

Evaluations of municipal AI tools, including 311 chatbots and vendor demonstrations, show recurring issues:

- Emergency-adjacent misdirection, such as suggesting rest and hydration for chest pain or shortness of breath instead of directing residents to 911.
- Different answers when residents imply different insurance status, neighborhoods, or language backgrounds.
- Fabricated eligibility rules for rental assistance or utility programs.

These patterns appear across systems currently marketed to cities and underscore why elected leaders and city managers should remain engaged with AI procurement decisions rather than delegating them entirely.

What City Leaders Should Ask

You don't need to run technical tests yourself. Your role is oversight. Before approving an AI tool, ask staff or vendors these five questions:

1. **Safety:** How does the system handle urgent or health-related questions? Does it consistently route residents to 911?
2. **Consistency and Fairness:** Will residents receive the same quality of service regardless of language, neighborhood, or how they phrase a question?
3. **Reliability:** Where does the AI get its information, and what happens when it doesn't know the answer?
4. **Transparency:** How will the city clearly inform residents they are interacting with AI, as TRAIGA requires?
5. **Oversight:** What documentation will the city keep to demonstrate due diligence?

What Responsible Adoption Looks Like

In November 2025, the City of Midland launched Jacky 2.0, an upgraded in-house AI chatbot supporting voice interaction, multiple languages, and resident-service navigation. The City scoped the tool narrowly, aligned it with service needs, and clearly communicated its capabilities and limitations.

The City of Kyle launched "Agent Kyle" to help answer questions and resolve issues reported through 311. Deployed in March 2025, the system has reported strong results, including faster service request processing and high first-call resolution rates. Kyle's approach included rigorous testing, employee education, and continuous fine-tuning before launch.

In March 2026, the City of Hutto launched HippoBOT, offering residents 24/7 access to city services in more than 70 languages.

Practical Next Steps

Responsible AI procurement does not require technical expertise. It requires knowing what good oversight looks like.

Ask vendors to prove their system works for your residents, not simply demonstrate it. That means documented testing using your community's real questions, languages, and service needs.

Insist on a pilot before full rollout. A 60- to 90-day supervised period with measurable benchmarks creates a meaningful decision point instead of a sales pitch.

Ask vendors what their system gets wrong. Every AI system fails sometimes. Vendors who understand and acknowledge failure modes know their products.

Before approving any AI tool, council should receive a concise summary answering the five oversight questions, outlining documented testing and known limitations, explaining the pilot scope and benchmarks, and detailing how the city will meet TRAIGA disclosure requirements. That single page can turn a sales pitch into a governable decision.

Moving Forward Under TRAIGA

Aligning a city's evaluation process with the NIST AI Risk Management Framework — a free, public standard — may also provide an affirmative defense under TRAIGA's safe harbor provisions.

TRAIGA does not block innovation but encourages responsible innovation. With clear expectations, strong documentation, and human oversight, Texas cities can improve resident services while protecting public trust. ★

Katryna Peart is a municipal AI governance researcher and consultant based in Leander. She is the creator of the Civic Pair Prompting evaluation method and has independently evaluated AI systems used by cities nationwide.