



# How Cities Can Communicate Clearly and Honestly About AI Use

Residents may have a range of perspectives about how artificial intelligence is used in local government. Concerns often stem less from the technology itself and more from uncertainty about how, when, and why it is being used. Proactive, plain-language communication can help build understanding and trust.

In many cities, AI tools are used for routine tasks like drafting meeting summaries or translating notices. However, without a clear explanation, some residents may assume broader uses such as surveillance, job reductions, or automated decision-making. Bridging that perception gap is primarily a communication challenge.

The cities handling this best are the ones that talk about AI the way they'd talk about any other tool: plainly, specifically, and before anyone has to ask.

## What Residents Actually Worry About

When cities survey residents or hold public forums about AI, the same concerns show up again and again. People want to know that someone is still in charge. They want to know their personal information isn't being fed into a system they don't understand. And they want the city to be upfront about what the technology is doing rather than burying it in a press release.

**Privacy** – Residents worry that their data, whether from utility accounts, permit applications, or 311 calls, will be loaded into public AI tools or shared with vendors who might use it for their own purposes. In Florida, many AI inputs and outputs related to city business may qualify as public records under state law. However, some data may still be protected by existing exemptions or confidentiality requirements under Chapter 119 when applicable. Telling residents this directly, in plain language on the city website, goes a long way toward calming fears before they take hold.

**Human oversight** – People are far more comfortable with AI when they know someone is still reviewing what it produces. That means saying explicitly: "AI tools assisted in preparing these materials, but all final content has been reviewed and approved by staff." When cities make clear that final decisions about permits, enforcement, or benefits stay with people, not software, the temperature drops quickly.

**Fairness and accuracy** – Most city AI use cases don't involve automated decisions at all. They involve drafting, summarizing, or translating, but residents don't know that unless someone tells them. A simple statement like "AI helps staff research and draft, but it doesn't approve or deny anything" will directly address the concern.

## How to Talk About AI

The most effective communication starts with the problem the city is solving, not the technology it chose. Consider the difference between two ways of describing the same thing. *"We're using AI to draft meeting transcripts so staff can publish them faster and residents can search them more easily."* This tells a resident everything they need to know. Compare that to *"We've adopted a generative AI solution to optimize workflows,"* which tells them nothing and may sound like it's hiding something.

Being specific about limits matters just as much. Residents trust cities more when leaders clearly state what AI does not do: it doesn't approve permits, it doesn't make enforcement decisions, it doesn't replace staff responses for complex issues, and every output gets reviewed before publication. Saying what a tool won't do is sometimes more reassuring than explaining what it will.

Plain language is non-negotiable. Most residents don't need to hear "large language model" or "machine learning pipeline." They care about outcomes: Is it faster? Is it accurate? Is someone checking it? Keep the explanation at that level unless the audience specifically asks for more.

And when something isn't perfect, say so. Acknowledging that AI makes mistakes, just like spell-check or auto-translation, builds more credibility than pretending the tool is flawless. "That's why our staff reviews all AI-assisted content before it's shared" is a complete answer to most concerns.

Cities should avoid hiding their use of AI. If residents learn about it through a news story or a social media post rather than from the city itself, the reaction can be worse. Even a perfectly reasonable use of AI can look suspicious when it surfaces by surprise.

## Two Florida Cities Doing This Well

**Marco Island** uses AI to assist with meeting documentation. The city publishes AI-generated meeting transcripts alongside the original source audio, clearly labels the tool used, and keeps staff responsible for official records and follow-up. The approach works because residents can see exactly what the AI produced, compare it to the audio themselves, and know that a person is still accountable for the final record.

**Tamarac** created a dedicated [AI page](#) on its website that lists every AI tool the city uses, by name, with a plain description of what each one does. Amazon Q Business handles internal knowledge queries for staff. Flock Safety runs the license plate recognition system for public safety. Microsoft Copilot and Power BI support internal dashboards that surface trends for city leadership. Each tool has a clear purpose, and the page explains the city's commitment to using AI in a way that comes across as transparent and secure.

Tamarac also joined the [Government AI Coalition](#) (GovAI Coalition), a national group of public agencies working on responsible AI use, led by the City of San Jose, California. That membership signals to residents that Tamarac isn't figuring this out alone, and it gives staff access to peer learning and shared standards across jurisdictions.

The lesson from both cities is the same: you don't need a perfect AI communication strategy. You need an honest one. Name the tool. Say what it does. Say what it doesn't do. Say who's accountable. Put it somewhere residents can find it.

## A Template Any City Can Use

For cities that want a starting point, here is sample language that can be adapted for a website, agenda packet, or public notice:

*"The City of [Name] uses AI tools to assist staff with tasks such as [meeting transcription / document drafting / translating public notices]. All AI-generated content is reviewed by city staff before publication. AI does not make decisions about permits, enforcement, or city services.*

*For a full list of AI tools the city uses, visit [link]. Questions? Contact [department/phone/email]."*

That kind of disclosure takes five minutes to write, costs nothing to post, and addresses the four things residents most want to know.

## Four Questions That Cover the Basics

Four key questions map a simple framework any city can use whenever it needs to explain AI use to residents, whether in a commission presentation, a web page, or a response to a public records request:

1. Why are we using AI?
2. What does the AI do, and what does it not do?
3. How is data protected?
4. Who is accountable?

Answer those four questions in plain language, and you've covered what most residents need to hear. Cities that want to go deeper, like Tamarac's full tool-by-tool disclosure page, can build on this foundation over time. Cities don't need to wait for a controversy to explain how they use AI.

The ones that communicate early and matter-of-factly are the ones that build trust before it gets tested. AI is just another tool in the municipal toolbox. What separates the cities that use it well from the ones that run into trouble isn't the technology. It's whether they talked to their residents about it first.