# FAQ: Before Your City Starts...

## Questions to Ask When Exploring AI in Local Government

This *Frequently Asked Questions* guide is designed to help Florida's municipal leaders ask the right questions before implementing AI tools. Each question is followed by a short explanation to guide your thinking. They highlight key considerations for risk, transparency, training, and alignment with state laws, such as the Sunshine Law.

The goal isn't to turn you into a technologist, but to help you approach AI adoption with confidence, accountability, and a clear focus on public value.

## Strategic Fit & Outcomes

**What public problem are we actually solving with AI?**
Be specific about pain points (e.g., permit backlogs, hurricane-season FAQs, or slow code-enforcement letters). If the problem isn't clear, the solution will drift, and residents may not feel the benefit.

**What does success look like in concrete metrics (time saved, cost avoided, service quality)?**
Pick a small set of measures you already track, such as days to permit, call center wait time, overtime costs, and satisfaction scores. Set a baseline now so you can show progress (not just activity) later.

**Is AI the best tool compared with simpler automation or process fixes?**
Map the workflow first. Many wins come from fixing forms, scripts, or handoffs. Choose AI when it clearly beats cheaper, lower-risk options.

**Who are the executive and department owners for this use case?**
Name one sponsor with authority (ex., city manager/deputy) and one operational owner (ex., department lead). Clear ownership speeds decisions, funding, and training.

## Governance & Policy

**Do we have a written AI use policy (employees, contractors, vendors)?**
Your policy should spell out approved tools, prohibited uses, privacy rules, and disclosure expectations. Make sure it aligns with Florida's public records and ethics requirements. Consider linking it to your computer acceptable use policy.

**Where should humans stay "in the loop," and who is accountable for reviewing outputs?**
Define which outputs always require human review (e.g., legal, finance, public-facing). Assign the approver by role (not a specific person) so it survives staff changes.

**How will the AI use policy be updated as laws and tools change?**
Keep it as an administrative directive or guidance you can refresh quarterly. Designate an owner (e.g., IT, City Attorney, Clerk) to track state and federal rules changes and make updates.

## Data Protection & Privacy

**What sensitive data will the AI tool need to ingest or access (PII, health, student, CJIS)?**
List data categories up front and decide what is never allowed in prompts. Favor the "minimum necessary" approach and redact where possible.

**Where does the data reside, and is the platform a government-grade environment?**
Confirm hosting (including U.S. data residency), encryption, and admin controls that fit government standards. Require single sign-on and audit logs so your Clerk and IT can trace usage.

**Will the AI tool vendor use our data to train their models, and can we opt out?**
Your contract should guarantee data isolation, clear retention limits, and deletion on request. If opt-out isn't available, treat that as a red flag.

## Bias & Accessibility

**How will we test for disparate impact and bias in outputs?**
Spot-check results across neighborhoods, languages, and demographics. Specify what "fair" means for the use case and measure against it regularly.

**Does the tool support plain language, multilingual, and culturally competent responses?**
Aim for sixth-to-eighth-grade readability and strong support for common languages in Florida, such as Spanish and Haitian Creole. Ensure content and chatbots meet accessibility standards (screen readers, captions) so everyone can use them.

## Accuracy, Quality, & Human Oversight

**Will outputs include citations with sources that can be reviewed manually to verify accuracy, such as when summarizing rules or laws?**
Require links back to ordinances, policies, or official web pages so staff can quickly verify. No citations means the content stays a draft until a human confirms it.

**Who reviews tone, legal accuracy, and policy alignment of drafted materials before release?**
Establish approvers by content type, such as the City Attorney for legal matters, Communications for tone, and the department lead for policy. Use a short checklist so reviews are fast and consistent.

**What is out-of-scope (no-go list) for this tool?**
Draw firm lines, such as no fully automated eligibility decisions, hiring recommendations, or law-enforcement determinations. A clear no-go list prevents quiet expansion into risky territory.

## Public Trust & Communications

**What is our public narrative for why and how we use AI?**
Connect AI use to better service (faster responses, clearer answers) and emphasize human accountability. Keep the message consistent across offices, departments, and the website.

**Do we require disclosure and/or labeling when AI assists in public-facing content?**
Consider whether to adopt simple labels like “AI-assisted, staff-reviewed,” and if you do, use them consistently. Disclosure builds trust and minimizes criticism later.

**What’s our crisis plan if an AI error goes public?**
Have a playbook: pause the AI tool, correct the record, notify affected residents, and document fixes. Identify the spokesperson and clearly outline the approval chain in advance.

**Have we taken steps to protect against "boilerplate voice" that could erode authenticity?**
Set a style guide and prioritize human editing of key messaging. Use local examples and stories so communications still sound like your city, not a template.

## Records Management & Sunshine/Open Records

**Are prompts, attachments, and outputs public records we must retain?**
Assume AI artifacts are public records under Florida's Sunshine Law. Coordinate with the Clerk (or other official custodian) on what to keep and for how long.

**How will we retain and search AI records?**
Pick a system of record (not someone's inbox) and capture prompts, outputs, dates, and approvers. Make it searchable so you can respond quickly to audits or requests.

**How will we handle redactions and exemptions for disclosed AI materials?**
Standardize redaction steps before release to avoid exposing PII or protected data. Train staff and spot-check compliance.

## Workforce, Training & Change Management

**Which roles will use AI tools, and how will jobs change (not just tasks)?**
Be transparent about what shifts (e.g., more review, less first-draft writing). Engage employees early to reduce anxiety and surface ideas.

**What is our training plan (prompting, verification, privacy, records)?**
Provide short, practical training with real city examples and clear "do/don't" rules. Refresh it periodically and require sign-off before access.

**Do we need an AI point person or team?**
Create a small cross-functional group (e.g., IT, Legal, Clerk, Communications) to approve tools, update policy, and handle questions. Treat them as an internal help desk for departments.

**What competencies will we build in-house versus rely on vendors for?**
Keep core skills – privacy, records, prompt discipline, quality review – inside the city. Use vendors for heavy build-from-scratch projects and integrations, or proprietary technologies, but require knowledge transfer where appropriate.