For Local Governments

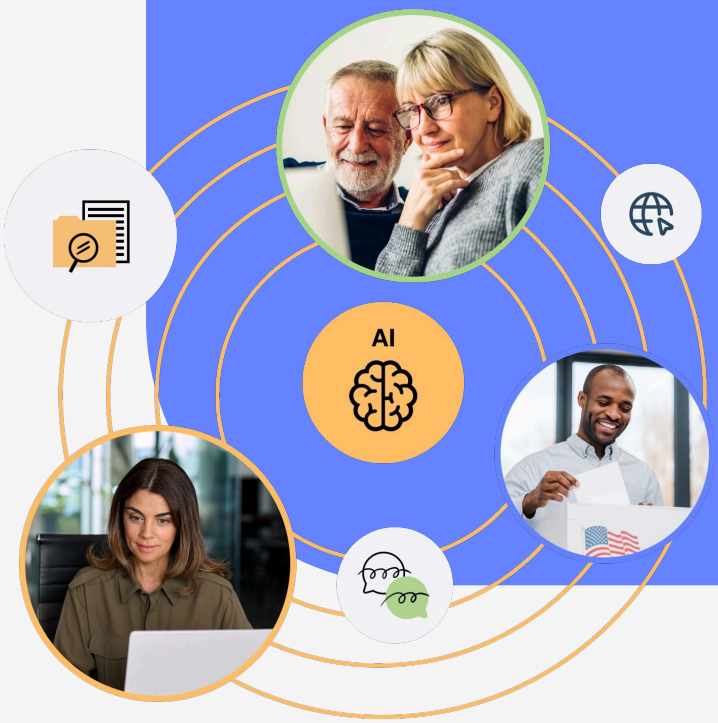

AI Readiness Toolkit

An actionable guide to help local governments safely and successfully adopt AI.

2026

Ordinal

Why Cities Need an AI Readiness Toolkit

Artificial Intelligence—especially Generative AI (GenAI)—has moved from the future into daily municipal operations.

Over 75% of county officials and staff already use GenAI tools both at work and in their personal lives according to the NACo AI Exploratory Committee Survey.

Local governments can no longer rely on avoiding AI; they must focus on safe, intentional, policy-aligned adoption.

This toolkit integrates the most respected guidance in the public sector (including NACo's AI County Compass, the OECD AI principles, the Michigan Municipal League recommendations, and the Virginia ODGA Data Readiness Checklist) and our experience as pioneers in government AI application to help cities move from hesitation to implementation.

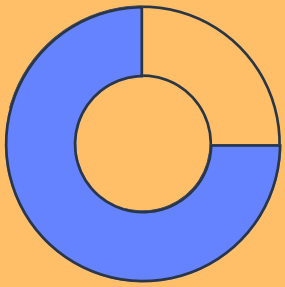

75%

already use GenAI tools both at work and in their personal lives

7 Readiness Pillars

We break out AI readiness into seven pillars:

Policy & Governance

Data Foundations

Responsible & Ethical AI Use

Workforce Preparedness

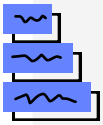

Prioritization

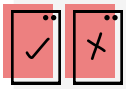

Choosing Vendors

Community Engagement & Transparency

These pillars will help you establish the policies and skills needed to effectively and safely roll out AI use across your municipality.

Pillar One

Policy & Governance

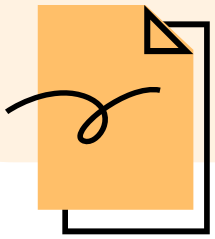

Build Your AI Governance Structure

Define Roles & Responsibilities

Who owns AI decisions, data, and outcomes?

AI risk increases when "everyone" is responsible—and no one actually is. Cities don't need a new department, but they do need clarity. Your structure will look different depending on your size, but here's a starting point:

AI Executive Sponsor

Typically:
City Manager, Assistant City Manager, or CIO

Responsibilities:
- Owns the why behind AI adoption
- Approves high-level AI use cases
- Resolves cross-department conflicts
- Communicates AI direction to Council

AI Governance Lead (Day-to-Day Owner)

Typically:
CIO, IT Director, or other Department leader with interest

Responsibilities:
- Maintains AI policies and standards
- Reviews proposed AI use cases
- Ensures compliance with privacy, security, and procurement rules
- Coordinates audits and reviews

Data Stewards (Per Dataset or System)

Typically:
Clerks, Records Managers, Department Heads

Responsibilities:
- Own specific datasets (e.g., ordinances, permits, resolutions)
- Approve data for AI use
- Ensure accuracy and updates
- Define retention rules

Legal / Privacy Advisor (Advisory Role)

Typically:
City Attorney or outside counsel

Responsibilities:
- Reviews high-risk use cases
- Advises on FOIA, public records, IP, and privacy
- Helps define disclosure language

Simple Operating Model

You can document all of this on one page with this template:

Role	Owner	Key Responsibilities
AI Sponsor	City Manager	AI strategy & accountability
AI Governance Lead	CIO	Policy, oversight
Data Steward	City Clerk	Data accuracy & approval
Legal Advisor	City Attorney	Risk & compliance

Important governance principle: All AI use within the city must include a designated **human in the loop**. Individuals using AI are accountable for its outputs, including accuracy, compliance, and appropriateness. AI systems do not make decisions on behalf of the city—people do. This responsibility should be clearly defined in AI governance documentation and reinforced through training. More details on this in the Workforce Preparedness Pillar.

Clear Data Ownership and Usage Rules

What data can be used, how, and by whom?

AI is only as reliable as the data behind it. Establishing clear rules for data ownership and usage helps ensure the right data is used in the right way by the right people—while protecting sensitive information and maintaining compliance.

Here's a quick risk level guide that can be helpful in structuring your data rules:

Public Data

Ex: Ordinances, agendas, resolutions, published codes & community information

Risk Level: 1

✅ Approved for all AI use (ideal for chatbots)

Internal Operational Data

Ex: Draft memos, internal emails, staff manuals

Risk Level: 2

⚠️ Approved for internal-only AI tools

Protected Data

Ex: PII, HR files, health data, some police data

Risk Level: 3

🚫 Prohibited in public or external AI tools, exceptions must be explicitly approved

Create Your City AI Policy

A clear, well-scoped AI policy is essential for responsible adoption. The goal of an AI policy is not to slow innovation, but to establish guardrails that protect residents, staff, and the city while enabling safe, effective use of AI tools.

At a minimum, your AI policy should clearly state that:

- Staff may not input protected or sensitive data into public or consumer AI tools, including personally identifiable information (PII), personnel records, health data, or confidential case information.
- Only approved datasets may be used to train or power city AI systems, and those datasets must have a clearly defined owner responsible for accuracy, updates, and oversight.
- All AI-generated outputs must be reviewed by a human before being shared publicly, relied upon for decision-making, or included in official communications or records.
- Verification and accuracy checks are mandatory—not optional. AI outputs should be treated as drafts or decision-support tools, not authoritative sources.

These requirements reinforce a core governance principle: AI does not replace human responsibility.

Here are some examples of AI policies that other cities have put together:

- **Seattle, WA →**
- **Tempe, AZ →**
- **San Jose, CA →**

Structuring Your AI Policy for Long-Term Relevance

To avoid policies becoming outdated as technology evolves, NACo recommends organizing AI governance using a policy pyramid:

- **Policies** define strict rules for what is permitted and prohibited. These are enforceable and change infrequently.
- **Standards** establish minimum expectations for how AI tools are evaluated, implemented, and monitored across the organization.
- **Guidelines** provide flexible, best-practice recommendations that help staff use AI effectively while allowing room for iteration as tools and use cases evolve.

This layered approach allows cities to maintain strong governance while remaining adaptable in a rapidly changing technology landscape.

Want some help drafting policy language?

Here's a template our team put together to get you started:

City AI Policy Template →

Pillar Two

Data Foundations

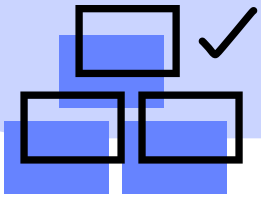

Assess Data Quality Before Implementing Any AI

A city's AI readiness is fundamentally about data readiness. Before deploying any AI tool, cities must clearly understand what data the system will use, where that data lives, and who is responsible for maintaining it over time. AI systems will amplify both strong data practices and weak ones, so taking the time upfront to assess your data quality is essential.

Cities should evaluate the following before implementation:

- **What data will the AI tool need access to?** Identify specific datasets required to support the use case, including structured and unstructured data (e.g., ordinances, permits, policies, FAQs). Avoid broad or open-ended data access.
- **What is the risk level of this data, and who owns it?** Classify the data as public, internal, or protected and assign a designated data owner responsible for accuracy, approvals, and oversight (see Pillar 1).
- **Is the data accurate, complete, and up to date?** Review data for gaps, outdated information, inconsistencies, or conflicting sources. AI tools rely on the quality of their inputs; poor data will produce unreliable outputs.
- **Where does the data live today?** Determine whether there is an existing centralized repository or source of truth, or whether data must be gathered from multiple systems, documents, or departments before use.
- **Who controls ongoing data additions, updates, and approvals?** Define who is responsible for approving new content, updating existing data, and retiring outdated information to ensure AI outputs remain accurate over time (see Pillar 1).
- **What is the process for identifying and resolving data issues?** Establish a clear workflow for reporting errors, reviewing data quality concerns, making corrections, and validating updates before they are reflected in AI tools.

Upgrade Cybersecurity As Needed Before Deploying AI

AI introduces new cybersecurity risks that cities must proactively address. Because AI systems can process, generate, and disseminate information at scale, existing security controls should be reviewed and strengthened before AI tools are deployed. Preparing your cybersecurity posture in advance helps protect city systems, staff, and residents while maintaining trust.

Cities should consider implementing or validating the following cybersecurity protocols:

- **Multi-factor authentication (MFA) and strengthened identity validation**
  Require MFA for access to AI tools, administrative consoles, and connected data systems to reduce the risk of unauthorized use or credential compromise.
- **Audit logs for AI tool usage**
  Ensure AI systems provide usage logging that captures who accessed the tool, when it was used, and what actions were taken. Audit logs support accountability, incident response, and compliance with records and oversight requirements.
- **Verification checks to identify AI-generated misinformation**
  Establish review processes and technical controls to detect inaccurate, misleading, or manipulated content generated by AI tools before it is shared internally or with the public.
- **Alignment with recognized security and compliance standards**
  Confirm that AI vendors meet relevant security and compliance requirements, such as SOC 2, HIPAA, CJIS, or other applicable standards, depending on the data involved and the department using the tool.

AI should be treated as part of the city's broader cybersecurity ecosystem—not as a standalone tool. Security expectations for AI systems should align with existing IT, data protection, and incident response policies, and be reviewed regularly as threats and technologies evolve.

Pillar Three

Responsible & Ethical AI Use

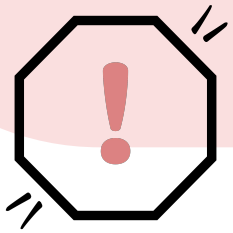

Adopt Core Ethical Principles

Responsible AI adoption begins with a clear ethical framework. NACo's Ethics Workgroup recommends that local governments ground all AI use in four core principles:

1. **Fairness, equitableness, and impartiality**
   AI systems should support fair and consistent outcomes and must not disproportionately disadvantage individuals or communities. Cities should ensure AI use does not reinforce existing inequities in access to services, information, or decision-making.

2. **Transparency**
   Cities should be open about when and how AI is used, particularly in resident-facing applications. AI systems should be explainable at a high level, and residents should be able to understand when they are interacting with AI rather than a human.

3. **Privacy**
   AI use must respect all applicable privacy laws and data protection requirements. Sensitive or protected information should be carefully controlled, and residents' data should never be used in ways they would not reasonably expect.

4. **Accountability**
   Humans—not AI systems—remain responsible for decisions, outcomes, and errors. Cities must maintain oversight, auditability, and clear escalation paths when AI-generated outputs are inaccurate or inappropriate.

Check For & Mitigate Bias

Bias is one of the most urgent and complex risks associated with AI. Bias can originate from multiple sources, including the design of the AI model itself or the historical and systemic biases embedded in the data used to train or inform it.

Use this checklist before AI-generated content is shared publicly, relied upon for decisions, or used in resident-facing services.

- **What decision or action could this output influence?**
  (If it affects access to services, information, or resources, apply heightened scrutiny.)
- **Is the AI using only approved, relevant, and up-to-date data?**
  (Outdated or partial data is a primary source of bias.)
- **Does the output treat all residents or users consistently?**
  (Would the response change unfairly based on who is asking or where they live?)
- **Is the language neutral, respectful, and free of assumptions?**
  (Watch for implied judgments about income, education, language ability, age, or background.)
- **Could this output unintentionally disadvantage or discourage any group?**
  (Even accurate information can create unequal outcomes if framed poorly.
- **Are important details missing that could mislead the user?**
  (Incomplete answers can introduce bias as much as incorrect ones.)
- **Can a human reviewer explain, at a high level, how this output was generated?**
  (If it can't be explained, it shouldn't be used.)
- **Has a human reviewed and validated this output for accuracy and context?**
  (AI outputs should never be published without review.)
- **Is there a clear way to correct or escalate if bias or errors are identified?**
  (Every AI output should have a correction path.)
- **Would you be comfortable defending this output publicly if questioned?**
  (If not, revise or withhold it.)

Pillar Four

Workforce Preparedness

Train Staff on the ‘Why’ and the ‘How’

AI training should go beyond tool demonstrations. Staff need to understand why the city is using AI, what problems it is intended to solve, and how their role fits into responsible use. This helps reduce misuse, builds confidence, and reinforces accountability.

At a minimum, training should cover:

- **Effective prompting**
  How to clearly and responsibly ask AI tools for information, drafts, or summaries—while avoiding leading questions, assumptions, or unintended bias.
- **Ethics and bias control**
  How bias can appear in AI outputs, how to recognize it, and how to apply the city’s bias review and human-in-the-loop requirements before using AI-generated content.
- **Privacy and data handling**
  Clear rules on what data may and may not be used with AI tools, including examples of protected data and common pitfalls that introduce risk.
- **When to escalate to a human**
  Guidance on when AI outputs are insufficient, unclear, or inappropriate—and when staff should pause, revise, or escalate to a supervisor or subject-matter expert.
- **How to verify AI outputs**
  Practical techniques for fact-checking, validating sources, and ensuring AI-generated information aligns with official city policies, codes, or records.

Cities may offer structured training such as Ordinal’s AI for Local Government Staff Workshop Series, which is designed specifically for municipal teams:

👉 https://www.ordinalforgov.com/ai-for-local-government-staff-workshop-series

Address Workforce Fears with Clarity

AI adoption can raise understandable concerns about job security, workload changes, and accountability. Cities should address these concerns directly and consistently.

AI should be positioned as:

- A support tool, not a replacement for staff
- A way to reduce repetitive or administrative work
- A means to help staff focus on higher-value, people-centered tasks
- A tool that still requires human judgment, oversight, and expertise

Clear communication helps build trust and prevents informal or unsafe use of AI tools driven by uncertainty or fear.

Develop a Multi-Year Workforce Capability Plan

AI readiness is not a one-time training event. Cities should plan for gradual capability-building over time, recognizing that tools, policies, and staff needs will evolve.

A sustainable workforce plan may include:

- **Role-specific training plans**
  Tailored guidance for different roles (e.g., clerks, planners, communications staff, front office teams) based on how AI is used in their daily work.
- **Partnerships with community colleges or universities**
  Collaborations to provide ongoing education, certifications, or workshops focused on digital literacy, data awareness, and responsible AI use.
- **Identifying AI champions within departments**
  Designating trusted staff members who can model good practices, support peers, and act as a bridge between departments and AI governance leads.
- **Including AI in new employee onboarding**
  Ensuring every new hire understands the city's AI policies, approved tools, data rules, and expectations from day one.

Pillar Five

Use Case Prioritization

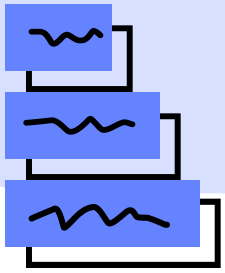

Start with Low-Risk, High-Impact Use Cases

Cities should begin their AI adoption journey with low-risk, high-impact use cases: applications that rely primarily on public or non-sensitive information and deliver immediate operational value.

These use cases allow staff to build familiarity with AI tools, establish governance practices, and demonstrate value without introducing unnecessary risk.

Risk Level: 1

Example Use Cases:

- **Agenda analysis and summarization**
- **Email summarization and drafting**
- **Public records Q&A**
- **Permit and process assistance**
- **Non-emergency call triage**
- **Public-facing chatbots built on official documents**

These use cases are ideal starting points because they:

- ✓ Rely on public or low-sensitivity data
- ✓ Do not automate decisions affecting resident rights
- ✓ Allow for clear human review and oversight
- ✓ Deliver immediate time savings for staff
- ✓ Improve access and responsiveness for residents

Be Cautious with Medium & High-Risk Use Cases

Not all AI use cases carry the same level of risk. As cities expand beyond basic applications, it is critical to differentiate between medium-risk and high-risk scenarios and apply stronger governance, oversight, and safeguards accordingly.

Risk Level: 2

Example Use Cases:

- **Policy drafting or internal guidance development**
- **HR tasks involving sensitive data**
- **Internal analytics involving partial PII**

Required safeguards for Medium Risk use cases:

- ✓ Use of approved, secure AI tools only
- ✓ Restricted access and role-based permissions
- ✓ Human review and validation of all outputs
- ✓ Clear documentation of data sources and usage

Risk Level: 3

Example Use Cases:

- **Criminal justice or law enforcement documents**
- **Health or human services data**
- **Automated recommendations or predictive analytics involving populations**

Required safeguards for Medium Risk use cases:

- ✓ Executive-level approval
- ✓ Legal and privacy review
- ✓ Strong data minimization and encryption
- ✓ Continuous monitoring and auditing
- ✓ Explicit prohibition on fully automated decision-making

Link Every Use Case to a Defined City Outcome

AI shouldn't be deployed simply because it is available or new. Every AI use case should be explicitly tied to a clear, measurable city outcome that aligns with operational priorities, service goals, or community expectations. This ensures AI investments are focused, defensible, and easier to evaluate over time.

When proposing or approving an AI use case, cities should be able to answer:
"What specific outcome will this improve, and how will we know?"

Below are common outcome categories with concrete examples.

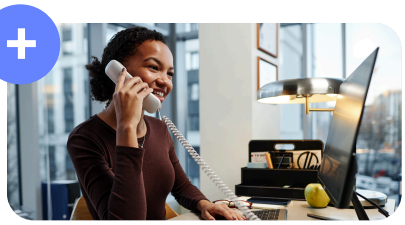

Reduce Service Backlog

What this looks like:
AI is used to help staff manage volume, not replace decision-making.

Example use cases:
- AI-assisted intake and categorization of 311 or front-office requests
- Summarizing large volumes of public records requests
- Drafting routine responses for common resident inquiries

How to measure impact:
- Reduction in average response time
- Fewer open or overdue requests
- Increased number of requests handled per staff member

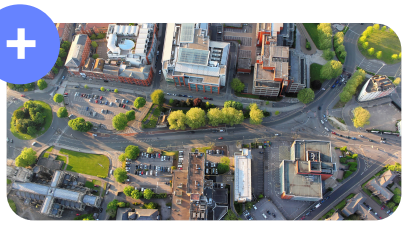

Improve Infrastructure Quality

What this looks like:
AI supports better planning, coordination, and insight.

Example use cases:
- Summarizing inspection reports or maintenance logs
- Identifying trends in service complaints related to infrastructure
- Supporting long-term planning with historical data analysis

How to measure impact:
- Improved prioritization of maintenance work
- Fewer repeat issues or service failures
- Better alignment between planning, operations, and budgeting

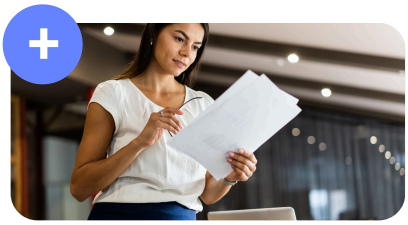

Decrease Staff Time Spent Searching for Information

What this looks like:
Staff spend less time hunting for documents and more time doing meaningful work.

Example use cases:
- Internal AI search across ordinances, resolutions, policies, and SOPs
- Agenda packet summarization for council or department meetings
- Rapid retrieval of historical decisions or precedents

How to measure impact:
- Time saved per task or request
- Reduced duplication of effort across departments
- Faster turnaround on internal requests

Pillar Six

Choosing the Right Vendors

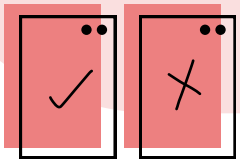

Asking the right questions and testing the AI functionality of vendors upfront is key to a successful implementation of AI into your city operations. Also important to note, many vendors are adding AI to their products, often silently. Cities must adapt quickly.

Questions to ask any AI vendor:

- Do you use our data to train external models?
- Do you comply with privacy laws (HIPAA, CJIS, state)
- Where is the data hosted? Is it in an approved, “government cloud” environment?
- Do you provide audit logs? How do we retrieve them for FOIA requests?
- How do you mitigate hallucinations?
- What model versions are used and how often updated?
- Do you use predictive modeling? What potential biases could this create?

Watch out for these red flags:

- "Our model is so smart it doesn't need to look up documents." (It will hallucinate).
- "It learns from your conversations to get smarter." (Privacy nightmare).
- "We can't cite sources because the AI synthesizes knowledge." (Unverifiable).
- “AI decisions are automated to save time.” (Removes the human in the loop.)
- “Our model is proprietary.” (Blackbox AI is not appropriate in government settings.)

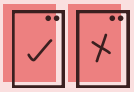

AI Procurement Clauses

As AI becomes embedded in more software products, cities should include standard AI-specific clauses in procurement contracts to protect data, ensure transparency, and maintain accountability. These clauses help prevent unintended data use and reduce risk as vendors continue to integrate AI into their platforms.

At a minimum, cities should consider including the following:

- **Data ownership and restrictions on reuse**
  The city retains full ownership of its data. Vendors may not reuse, sell, or repurpose city data beyond providing contracted services.
- **Retention and deletion requirements**
  Vendors must clearly define how long data is retained and provide mechanisms for secure deletion upon request or contract termination.
- **Mandatory AI disclosures**
  Vendors must disclose where and how AI is used within the product, including any changes to AI functionality over time.
- **Accessibility compliance**
  AI-powered features must comply with applicable accessibility standards (e.g., WCAG) to ensure equitable access for all residents and staff.
- **Minimum security measures**
  Vendors must meet baseline security requirements appropriate to the data involved (e.g., SOC 2, encryption, access controls, audit logging).
- **Prohibition on training external models with city data**
  City data may not be used to train, fine-tune, or improve external or third-party AI models without explicit written approval.

Because AI capabilities are increasingly added to products by default, cities should embed AI terms into all procurement processes moving forward, rather than treating AI as a special or one-off exception. This ensures consistent governance as technology evolves.

Pillar Seven

Community Engagement & Transparency

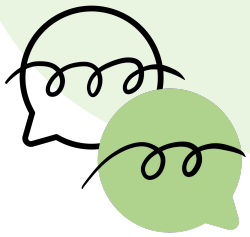

Public trust is essential to responsible AI adoption. Cities that communicate early and clearly about how AI is used (and how it is not used) are better positioned to address concerns, reduce misinformation, and build confidence among residents and stakeholders.

Build Public Trust Through Open Communication

Research from NACo highlights several common misconceptions about AI in local government, including the belief that cities can avoid AI entirely by banning it, or that AI systems are always accurate. Left unaddressed, these myths can undermine trust and fuel unnecessary concern.

Cities should proactively communicate:

- *How AI can help residents*, such as improving access to information, reducing wait times, or making services easier to navigate.
- *What safeguards are in place,* including human review, data protections, privacy controls, and limitations on how AI can be used.
- *When residents are interacting with AI*, particularly in public-facing tools like chatbots or virtual assistants, with clear disclosure and a path to reach a human when needed.

Clear, consistent messaging helps residents understand that AI is a tool to support service delivery, not a replacement for accountability or human judgment.

Resident-Facing Benefits to Highlight

When discussing AI with the public, cities should focus on tangible benefits that align with community needs and values. Common benefits to highlight include:

- **Improved service delivery**
  Faster answers, fewer delays, and quicker access to commonly requested information.
- **Increased equity and accessibility**
  Multilingual support, plain-language explanations, and expanded access to services outside normal business hours.
- **Personalized and proactive support**
  Better guidance based on resident questions or needs, without requiring residents to navigate complex systems on their own.
- **Optimized social services**
  Improved coordination and earlier identification of service needs, while maintaining human oversight and privacy protections.
- **Smarter resource allocation**
  Better forecasting and planning based on trends and historical data, supporting more informed decision-making by city leaders.

Not sure where to start?

Here's a template our team put together to explain important aspects of AI use in your city.

AI FAQ Template→

Bonus

AI Maturity Model for Cities

This model helps a city assess its current state of AI readiness and identify practical next steps. Cities do not need to reach the highest level immediately—progression is incremental and should align with capacity, risk tolerance, and community needs.

Level 1: Awareness

1

AI use is informal and uncoordinated.

- Staff may be experimenting with public AI tools on their own
- No formal AI policies, governance, or guidance
- Data ownership and usage rules are unclear or undefined
- Leadership awareness is emerging, but strategy is not yet established

🚨 **Primary risk:** Inconsistent use, privacy exposure, and lack of accountability

💡 **Primary opportunity:** Establish baseline policies and education

Level 2: Foundational

2

Governance and guardrails are being established.

- Core AI policies and acceptable use guidelines are in place
- Data classification and ownership are defined
- Leadership has acknowledged AI as a strategic issue
- Initial vendor and procurement considerations include AI

🚨 **Primary risk:** Overly cautious adoption or stalled momentum

💡 **Primary opportunity:** Move from policy to practice

Bonus

AI Maturity Model for Cities

Level 3: Early Adoption

3

AI is delivering value in controlled, low-risk ways.

- Approved low-risk use cases are live (e.g., document search, summarization)
- Staff training programs are underway
- Human-in-the-loop review is consistently applied
- Early metrics are being tracked (time saved, usage, access)

🚨 **Primary risk:** Siloed implementations

💡 **Primary opportunity:** Expand thoughtfully across departments

Level 4: Integrated

4

AI is embedded into daily operations with oversight.

- AI tools are used across multiple departments
- Clear audit trails, approval workflows, and governance processes exist
- Data updates and quality management are routine
- Outcomes and impact are measured and reported

🚨 **Primary risk:** Scaling faster than governance evolves

💡 **Primary opportunity:** Optimize workflows and resident experience

Level 5: AI-Optimized

5

AI is a trusted, transparent part of city operations.

- Predictive analytics support planning and resource allocation
- Automated workflows handle routine tasks with human oversight
- Residents have reliable 24/7 self-service access to information
- High transparency around AI use builds public trust
- Continuous improvement is built into governance and training

🚨 **Primary risk:** Over-reliance without ongoing review

💡 **Primary opportunity:** Lead by example in responsible innovation

Ordinal

Ready to bring AI into your city?

Speed up your teams efficiency with government-safe AI tooling. Research your internal documents instantly, draft staff reports in minutes, or implement an AI chatbot for residents.

Learn More →

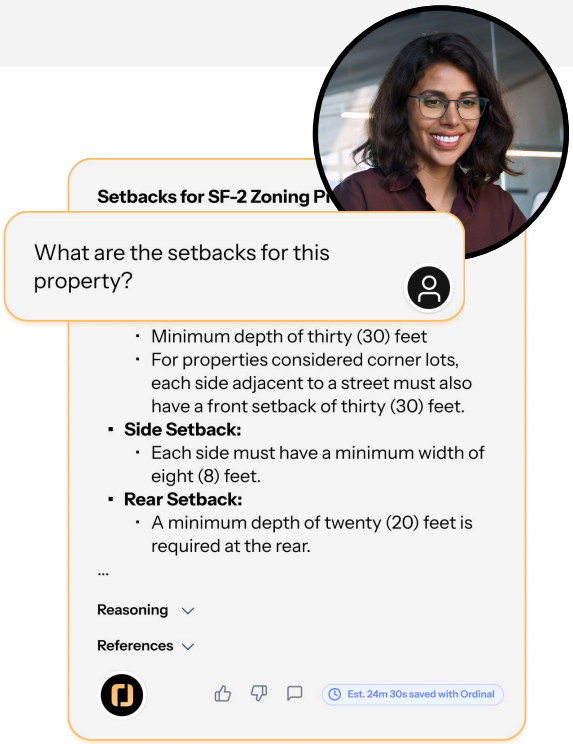

About Ordinal

Ordinal is an AI-powered research assistant that lets your staff find and share consistent, accurate answers in seconds.