2025

AURORA ILLINOIS

# AI Roadmap

TURNING STRATEGY INTO ACTION THROUGH RESPONSIBLE, PHASED AI ADOPTION

KARUMURI, LEELA

# Table of Contents

# Change Log

This document should be reviewed and updated at least annually or whenever significant changes in the technical, city governance, or operating environment warrant.

| | |
|---|---|
| Document Owner | City of Aurora Chief Information Officer (CIO) |

| Version No. | Modified on Date | Modified By | Change Description | Change Impact |
|---|---|---|---|---|
| 1.0 | October, 2025 | IDC | Drafted roadmap | Major |
| 1.1 | Nov, 2025 | COA | Minor | Minor |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |

# Executive Summary

**Phase 1: FOUNDATIONS (0–12 Months)**

*Building the ground on which AI can stand*

| # | Initiative | Target Completion (Month) |
|---|---|---|
| **1** | Establish AI Governance Committee | 3 |
| **2** | Develop Foundational AI and Data Policies | 6 |
| **3** | Build Centralized AI Resource Hub | 6 |
| **4** | Develop AI-Specific Vendor and Procurement Framework | 6–7 |
| **5** | Conduct Data Survey and Quality Assessment | 9 |
| **6** | Launch Workforce Skills Assessment | 9 |
| **7** | Formalize Change Management and Communication Plan | 10 (ongoing) |
| **8** | Initiate Low-Risk AI Pilots | 12 |

**Phase 2: INTEGRATION (12–24 Months)**

*Making AI part of the machinery*

| # | Initiative | Target Completion (Month) |
|---|---|---|
| **1** | Integrate AI Procurement Standards into Citywide Contracts | 15–16 |
| **2** | Introduce Role-Based AI Certification Programs | 15 |
| **3** | Launch Resident Transparency Tools | 16–17 |
| **4** | Embed AI into IT Architecture and Standards | 17 |
| **5** | Deploy AI-Ready Cybersecurity Controls | 19 |
| **6** | Expand Data Quality and Integration Efforts | 19–20 |
| **7** | Reengineer High-Volume Workflows | 20 |
| **8** | Scale Successful Pilots to Production | 18–22 |
| **9** | Implement Knowledge Management Systems | 24 |

**Phase 3: OPTIMIZATION (24–36 Months)**

*Institutionalizing AI for continuous improvement and innovation*

| # | Initiative | Target Completion (Month) |
|---|---|---|
| **1** | Integrate Structured and Unstructured Data | 28 |
| **2** | Operationalize Knowledge Management | 32 |
| **3** | Scale Transformative AI Projects | 33 |
| **4** | Expand Civic Partnerships and Innovation Ecosystem | 34 |
| **5** | Institutionalize AI Oversight and Continuous Improvement | 36 |
| **6** | Enhance AI Resilience and Compliance | 36 |

# Background

The strategy established the City's overarching principles, vision, and strategic objectives for responsible artificial intelligence (AI) adoption across all municipal functions, constituents, stakeholders, and partners. It defined the 5C Principles — Confidence, Confidentiality, Contribution, Conscience, and Clarity — as the ethical and operational foundation for all AI activity within the City's government.

The AI Strategy also identified key opportunity areas, current capability gaps, and a phased "Walk–Jog–Run" maturity model to guide adoption. The document provided a directional view of *why* and *where* AI should be deployed, emphasizing the need for structured governance, robust data management, and workforce preparedness before pursuing high-impact use cases.

Building on that foundation, this AI Roadmap translates the strategic framework into an actionable, time-bound plan spanning three years (2026–2029). It provides:

- **An ordering** of initiatives showing which actions must occur in order and which may proceed in parallel.
- **Defined dependencies and milestones** that connect governance, data readiness, workforce development, and constituent engagement activities.
- **Measurable events** that demonstrate progress toward the City's AI vision and its practical realization.

Together, the AI Strategy and this AI Roadmap form a comprehensive framework for responsible AI adoption for the City; the strategy defines intent and principles, while the roadmap defines execution and sequencing. Both documents should be treated as living instruments, reviewed annually to reflect changing technical, legal, and organizational conditions within the City and the broader public sector environment.

Methodology and Design Principles

The roadmap applies a "Walk–Jog–Run" methodology across a 36-month period.

This phased model is designed to:

- Minimize risk by proving value early through low-risk pilots.
- Build organizational and stakeholder confidence and capacity before scaling.
- Embed sustainability by aligning data, governance, and workforce readiness from the outset.
- Ensure transparency and equity in all initiatives.

Each phase has distinct objectives but shared enablers, e.g., data quality, workforce training, and governance, which evolve in parallel as the City delivers added value.

## Summary of the Three Phases

| Phase | Timeline | Theme | Primary Focus | Value Outcomes |
|---|---|---|---|---|
| **Phase 1: Foundations (Walk)** | 0–12 months | Building Trust and Structure | Establishing governance, policies, and data readiness; launching pilot projects | Confidence in results; organizational readiness; early efficiency wins |
| **Phase 2: Integration (Jog)** | 12–24 months | Scaling Value Across Functions | Embedding AI within IT architecture and city workflows | Contribution through measurable efficiency and service improvements |
| **Phase 3: Optimization (Run)** | 24–36 months | Sustaining Innovation and Resilience | Institutionalizing governance, scaling transformative use cases | Adaptive AI ecosystem |

# Maintenance

This roadmap is designed with a three-year horizon and should be revisited annually to reflect changing economic, political, technical, resourcing, and budgetary conditions. This ensures the City's AI adoption remains relevant, sustainable, and responsive to community needs.

# Phase 1: FOUNDATIONS (0–12 Months)

### "Building the Ground on Which AI Can Stand"

The goal is to ensure that all future AI work operates within a secure, transparent, and ethical framework.

## 1. Establish the AI Governance Committee

The City will formally establish an AI Governance Committee to provide cross-departmental oversight, coordination, and accountability for all AI-related initiatives. The committee will bring together leaders from IT, Legal, HR, Procurement, Communications, and Community Engagement (where applicable), ensuring that both technical and public-facing considerations are addressed in every stage of AI planning and implementation. This diverse composition reflects the City's commitment to balancing innovation with ethics, risk management, and transparency.

The committee's mandate is to define roles, oversight procedures, and risk thresholds for AI adoption. It clarifies what types of projects require formal review, what standards apply to data use, and how compliance with City policies will be maintained. It will also oversee the intake process for proposals, helping prioritize efforts that deliver public value while managing risk.

The outcome of this effort will be a concise AI Governance Charter and Operating Model (Month 3) that codifies responsibilities, decision rights, and reporting mechanisms, establishing a durable structure for responsible AI governance.

**Deliverable:** AI Governance Charter and Operating Model (Month 3)

## 2. Develop Foundational AI and Data Policies

The City will establish a foundational policy framework to ensure that all AI initiatives are guided by clear principles of ethics, confidentiality, transparency, and accountability. This work builds on existing policy progress, for example, the Data Policy and AI Governance

Policy, which are either already approved or in advanced stages of development. These core documents serve as anchor references, ensuring that any new guidance aligns with the City's broader commitments to privacy, security, and responsible technology use.

Additional supporting policies will address key operational areas such as vendor compliance, explainability requirements, and fair use of data in AI systems. Together, these policies will form a cohesive AI Policy Suite (Month 6) that defines how AI may be proposed, developed, and deployed across City departments. The focus is on establishing guardrails early so pilot projects, procurement decisions, and staff training occur within a shared framework of ethical practice and data integrity, reinforcing public trust from the outset.

**Deliverable:** AI Policy Suite (Month 3)

## 3. Develop AI-Specific Vendor and Procurement Framework

As AI capabilities are available as both standalone tools and imbedded in licensed software, the City will need a clear but lightweight framework to evaluate, procure, and manage these solutions responsibly.
This is not about creating an entirely new procurement system but rather adding AI-specific guardrails and questions to the processes already in place.

The goal in this first phase is to ensure that procurement and vendor management teams know what to look for and what to ask when AI-related products or services are proposed. The framework will cover a small set of core principles:

- **Transparency:** Vendors must disclose whether and how AI is used in their products, including model provenance, training data sources, and any embedded generative components.
- **Data Ownership and Use:** Contracts should clarify who owns the data and outputs, how data may be stored or reused, and how privacy is protected.
- **Explainability and Bias Mitigation:** Vendors must demonstrate reasonable steps to ensure that AI-driven recommendations or decisions can be explained, and that fairness or bias concerns are monitored.
- **Security and Compliance:** Products must align with City privacy policies, cybersecurity controls, and applicable state or federal requirements.

This framework will take the form of a concise AI Vendor Management Handbook or Checklist (Month 6-7). It is essentially a short reference guide or checklist that is appended to RFPs and contract templates.

The emphasis is on practical adoption, not red tape. Procurement officers and project leads can use it to ask the right questions, document AI-specific risks, and ensure vendor accountability without delaying normal purchasing cycles. As the City gains experience,

this initial handbook can be refined later to include scoring criteria or model-evaluation templates for higher-risk AI procurements.

**Deliverable:** AI Vendor Management Handbook or Checklist (Month 6-7)

## 4. Build Centralized AI Resource Hub

The City will stand up a centralized online AI Center of Enablement to serve as the single point of access for all information, tools, and guidance related to AI adoption. This site will house (or link to) essential resources such as approved policy documents, governance templates, training materials, approved tools with their use-cases, and FAQs, along with updates on active AI projects and pilots. It will also include quick reference materials to help departments understand the City's AI principles, responsible-use expectations, and links to vendor and data management standards.

To make adoption straightforward and accountable, the Center will also provide clear intake and review processes for AI proposals. Simple digital forms will guide requesters through a short checklist that includes purpose, data used, expected outcomes, and alignment with City policies, so proposals can be reviewed consistently by the AI Governance Committee. This process will help the City track new ideas, ensure appropriate oversight, and build a transparent record of AI activity without adding procedural complexity.

**Deliverable:** AI Center of Enablement (Month 6)

## 5. Conduct Data Survey, and Targeted Inventory and Quality Assessment

Rather than attempting a full enterprise-wide data inventory, which would be impractical given the City's volume and diversity of data and data silos, this activity focuses on identifying high-value, high-risk data domains that will most directly influence the success and integrity of early AI initiatives.

The goal is to understand what data exists, where it resides, and what risks it may introduce when used in AI contexts, not to achieve exhaustive coverage. Departments will collaborate with the City's Data Governance team to perform a tiered data survey emphasizing:

- **High-value datasets** that are foundational for initial AI use cases (e.g., constituent service data, permitting, 311 requests, inspection records, or public works scheduling).
- **High-risk or sensitive data sources**, including any information containing personally identifiable, quasi-identifiable, or derived data that may present privacy, bias, or ethical risks in AI applications.

- **Quality and accessibility indicators**, assessing completeness, structure, duplication, and ownership clarity within those high-value and high-risk areas.
- **Documentation of lineage and usage constraints**, ensuring that departments understand how existing data may or may not be safely applied in AI projects. Ensure identity and access privileges are supported in any AI use.

The outcome of this activity will be a Targeted Data Inventory and Risk Assessment Report, providing:

- A prioritized map of key data assets relevant to near-term AI efforts.
- Identification of potential compliance, privacy, or data integrity risks.
- Recommendations for improving data governance and quality in those critical domains.

This measured, risk-based approach avoids "boiling the ocean" while establishing the confidence and situational awareness necessary for responsible AI adoption and future data platform investments in Phase 2.

**Deliverable:** Targeted Data Inventory and Risk Assessment Report (Month 9)

## 6. Launch Workforce Skills Assessment

The City will conduct a light-touch assessment to understand its current workforce readiness for AI. It is not a major HR initiative, but a focused check-in to see where strengths and gaps exist.

The effort will be under the direction of the AI Governance Committee with Human Resources input, in collaboration with leadership using short surveys and interviews with key staff. The goal is to quickly determine:

- **Current familiarity** with data, analytics, and basic automation tools that could be leveraged in early AI projects.
- **Critical readiness gaps** in areas such as AI ethics awareness, data-governance practices, and the responsible use of generative tools.
- **Immediate upskilling needs** for the first wave of pilot users and project sponsors (for example, staff likely to use or oversee AI in constituent services or communications) and technology staff.

The assessment should take weeks (not months) and focus on practical insight rather than exhaustive documentation. Its output, the AI Readiness and Skills Gap Snapshot (Month 9),

will summarize where the City can act quickly (e.g., short training modules, peer mentoring, vendor workshops) and where deeper development may be needed later.

This targeted approach provides enough visibility to guide early training investments without burdening staff or delaying other foundational work.

**Deliverable:** AI Readiness and Skills Gap Snapshot (Month 9)

## 7. Formalize Change Management and Resident Communication Plan

Introducing AI into City operations represents more than a technical shift; it changes how employees work, how residents interact with services, and how the City explains decisions to the community.

To ensure this transformation is constructive and trusted, the City will establish an intentional approach to change management and communication early in its AI journey.

This effort begins by defining clear, consistent messages that explain what AI is, and just as importantly, what it is not. Communications will emphasize that AI is a tool to support staff and improve service quality, not replacing human judgment. Departments will coordinate with the Communications Office, HR, and union representatives to align messaging with workforce realities and address potential concerns proactively.

Key actions will include:

- Developing early communications and FAQs for both employees and residents, outlining the City's principles (transparency, ethics, privacy) and describing the specific benefits of upcoming pilot projects.
- Integrating AI messaging into existing internal channels such as newsletters, department briefings, or leadership updates to normalize discussion of AI and reduce uncertainty.
- Creating a basic resident-facing information page or microsite explaining how AI is used in City operations, how personal data is protected, and where citizens can ask questions or provide feedback.
- Engaging labor and HR partners to ensure that AI adoption respects existing agreements, supports staff development, and reinforces confidence in fair labor practices.

This plan will lay the groundwork for trust, transparency, and staff participation, ensuring that AI initiatives are introduced with clarity and empathy.

The resulting AI Change Management and Communication Framework will become the foundation for ongoing workforce engagement and public outreach as AI adoption expands in later phases

**Deliverable:** AI Change Management and Communication Framework (Month 10, ongoing)

## 8. Initiate Low-Risk AI Pilots

Early pilots will serve as proof-of-value experiments to help the City build familiarity, confidence, and internal know-how with AI before scaling to higher-risk or public-facing initiatives.

These pilots are intended to be lightweight, fast, and reversible. They should be testing ideas in controlled environments where outcomes can be observed and measured without creating significant operational or reputational exposure.

The City should begin with three to four focused projects that:

- Address clearly defined, narrow challenges (for example, document summarization for internal reports, a chatbot answering resident FAQs, or RFP creation or vendor response summarization).
- Use existing, non-sensitive data where possible to minimize privacy and security risk.
- Delivery results in weeks rather than months, with an emphasis on learning rather than perfection.
- Involve staff directly so that pilots double as hands-on training for teams that will later guide broader adoption.

Each pilot will have a simple "learning plan", i.e., what the City wants to understand about AI performance, usability, and governance implications (in addition to use case benefits). Outcomes will feed into a short Lessons-Learned Summary, highlighting:

- Where AI created measurable value (time savings, accuracy, satisfaction).
- Where risks, quality gaps, or governance challenges emerged.
- What changes should be made to City AI standards, data policies, or training materials before scaling.

The aim is to start small, learn fast, and adapt. These early wins will generate internal enthusiasm, refine oversight processes, and create practical examples to guide larger, integrated efforts in Phase 2.

**Deliverable:** Pilots (Month 12)

**Parallel Workstreams**

- **AI Literacy and Ethics Training:** Begin baseline training for all staff.

- **Resident Outreach:** Launch awareness campaign emphasizing transparency, accessibility, and privacy safeguards.
- **Partnership Development:** Engage universities, peer municipalities, and vendors for co-learning and joint pilots.

**Expected Outcomes**

By the end of Year 1:

- AI use is governed by clear rules, policies, and oversight.
- Procurement guardrails are in place to guide early vendor engagements and pilot selection.
- Staff and residents have foundational awareness of AI principles.
- Pilot projects demonstrate quick wins and yield measurable efficiency gains.
- The City has an accurate picture of its data assets and workforce readiness.

# Phase 2: INTEGRATION (12-24 Months)

"Making AI Part of the Machinery"

The second phase transitions from planning to scaling and integration, focusing on embedding AI into the City's operational and technical fabric. This isn't about sweeping transformation; it's about smart, deliberate integration and making AI a dependable part of the City's systems, staff, and residents interactions.

The goal is to drive early successes into repeatable capabilities where integrating AI within IT architecture, daily workflows, and constituent-facing services delivers measurable, sustainable public value.

Phase 2 is where AI begins to shift from pilot to operational improvement. Systems become interconnected, governance becomes part of normal decision-making, and data management, cybersecurity, and staff enablement evolve to support AI at scale.

**Key Sequential Initiatives**

## 1. Integrate AI Procurement Standards into Citywide Contracts and Vendor Reviews

With the AI Vendor Management Handbook established in Phase 1, the City will move in Phase 2 to embed those standards into everyday procurement and vendor oversight. The focus will be on normalizing AI due diligence as part of existing purchasing and contract

renewal processes, rather than creating new layers of review. Procurement officers, contract managers, and department leads will begin to apply AI-specific questions and guardrails consistently across all technology procurements and renewals.

Key activities will include:

- Updating standard RFP and contract templates to incorporate AI disclosure language, data ownership terms, and explainability expectations.
- Integrating AI evaluation criteria into vendor scoring models, especially for procurements that involve automation, analytics, or decision-support tools.
- Reviewing existing vendor agreements where AI functionality has been added post-purchase to ensure continued compliance with City policies.
- Engaging with Legal and IT security to refine standard clauses around privacy, data reuse, and ethical AI practices.
- Creating a light training module for procurement staff to ensure consistent application of AI guidelines.

This step ensures that AI-related risks are considered early in the process, vendors remain transparent and accountable, and the City's ethical and security standards are upheld without slowing operations.

**Deliverable:** Updated Procurement Templates and AI Contract Review Guidelines (Month 15-16)

## 2. Introduce Role-Based AI Certification Programs

As the use of AI expands, staff need more than awareness; they need the confidence and practical skills to apply AI responsibly in their daily work. Phase 2 will introduce role-based certification programs tailored to different groups across the organization. Analysts, planners, and communicators will receive focused training on how to use AI tools effectively, interpret AI outputs, and manage them within City policies.

To accelerate development, the City may identify third-party partners or content providers experienced in AI training and compliance education, similar to the model used with KnowBe4 for security awareness programs. These partners can assist in creating modular, accessible content that reflects both public-sector realities and evolving AI best practices. The first cohort of certified participants will complete training by Month 15, aligning with the rollout of scaled AI tools. Graduates will serve as departmental AI champions, reinforcing consistent and responsible adoption across the organization.

**Deliverable:** Role-Based AI Certification Program, Partner Engagement, and First Cohort Completion (Month 15)

## 3. Launch Resident Transparency Tools

As AI projects move from pilot to production, communication will evolve from initial awareness toward active transparency and two-way engagement. The City will build on its foundational change management and messaging work to make AI use visible, understandable, and open to feedback.

Key activities include:

- **Expanding resident communication channels** to share updates on AI-enabled services, performance improvements, and safeguards.
- **Launching public transparency dashboards** showing how AI is used, what data it draws from, and how fairness, privacy, and quality are monitored.
- **Establishing feedback loops**, such as surveys, community sessions, and online forms, to capture staff and resident sentiment and inform ongoing governance and training improvements.
- **Continuing workforce engagement**, reinforcing that AI is an assistive capability designed to improve efficiency and public service outcomes, while maintaining respect for roles and agreements.

By the end of this phase, the City's communication and engagement model will have shifted from explaining AI to demonstrating accountability and responsiveness, and turning communication from a one-time activity into a consistent mechanism for trust and continuous improvement.

**Deliverable:** More Robust Transparency Tools and Processes (Month 16-17)

## 4. Embed AI into IT Architecture and Standards (AI Architectural Blueprint)

The City will integrate AI into its core technical and architectural foundation, establishing the frameworks needed for interoperability, scalability, and security.
This includes:

- Establish interoperability and API standards so AI systems can interoperate cleanly with existing platforms.
- Build AI governance and risk checks into the enterprise architecture review process, so new projects automatically pass through the checkpoints.
- Coordinate between IT, cybersecurity, and data teams to ensure AI projects align with modernization efforts already underway.

By mid-year, these standards will come together as an AI Architectural Blueprint, providing a consistent foundation for new deployments without stifling experimentation.

The outcome will be a repeatable architecture for future AI deployments that makes new projects faster to launch and easier to manage.

**Deliverable:** AI Architectural Blueprint (Month 17)

## 5. Scale Successful Pilots to Production

Building on the lessons from Phase 1, the City will expand only those pilot projects that have shown clear, measurable value and can be operated safely at larger scale.
The focus will be on simple, pragmatic scaling, i.e., turning a few well-tested use cases into sustainable tools that teams can rely on day-to-day.

Examples might include expanding the resident-service chatbot to additional departments, deploying the document-summarization assistant across key administrative units, or integrating automated triage into existing 311 and permitting workflows.

Before scaling, each pilot will undergo a governance review to confirm that:

- The data used is appropriate and securely managed.
- Results meet accuracy and fairness expectations.
- Staff are comfortable using and maintaining the tool.

Departments will then transition pilots into production through existing IT or vendor channels, using the same playbook developed in Phase 1. This keeps implementation manageable, repeatable, and low-risk, while expanding the tangible benefits of AI across City operations.

The emphasis remains on value, not volume. Scaling only where success is proven and operational readiness exists.

**Deliverable:** Successful pilots to production (Month 18-22)

## 6. Deploy AI-Ready Cybersecurity Controls

As AI becomes a core capability, cybersecurity needs to evolve alongside it. AI will play a dual role in the City's security posture as both a tool for defense and a new category of assets requiring protection. It will also be a threat to defend against.

Efforts will include:

- Implementing zero-trust architectures to secure data and systems against unauthorized access.
- Deploying AI-assisted detection and response tools to augment analyst capacity, reduce fatigue, and accelerate incident response.

- Updating cybersecurity playbooks to include AI threat scenarios such as model manipulation or prompt injection. And direct use of AI by malicious actors.
- Training analysts and administrators to interpret AI-driven security alerts and manage automation responsibly.

This step reinforces that AI adoption doesn't just improve productivity — it also enhances the City's ability to defend itself in an increasingly complex digital landscape.

**Deliverable**: Enhanced AI-Ready Cybersecurity Controls (Month 19)

## 7. Expand Data Quality and Integration Efforts

If Phase 1 helped the City understand what data it has and where it lives, Phase 2 is about making that data work together. AI can only be as smart as the information it draws from, and right now, the City's data sits in silos (with varied levels of reliability and trust across the city). It is difficult to combine or use systemically. The task now isn't collecting more; it's to connect what already exists and make it useful, safe, and trustworthy.

This phase expands the City's data foundation in several key ways:

- **Unifying access across departments:** Build bridges between systems that store related information, for example, connecting 311 requests to Public Works schedules, or inspection data to permitting workflows, so departments can see a fuller picture.
- **Extending governance to new data types:** Move beyond structured data to include text, PDFs, images, and other unstructured materials. This doesn't mean indexing everything; it means tagging and managing what has value for analytics or AI.
- **Creating a centralized platform for discovery:** Develop a **data lake** that consolidates high-value datasets while maintaining strong privacy, security, and access controls. Departments can contribute data safely and retrieve it when needed, without losing ownership.
- **Embedding privacy and quality checks:** Apply consistent metadata standards, version tracking, and simple quality metrics (e.g., completeness, duplication, age) to ensure the data feeding AI systems is dependable.

The result will not be a perfect or all-encompassing data warehouse. That is neither realistic nor necessary. Instead, the City will have a shared, governed foundation that supports both immediate AI projects and future innovation. Start with the most recent data and work backwards. Don't try to address it all; that is both too time and resource intensive (and often unnecessary).

As this work matures, staff will begin to see the practical benefits: reports and dashboards pulling from a single source of truth, fewer redundant datasets floating across departments, reduced manual work, and cleaner, better-structured inputs for AI efforts. These incremental improvements make the City's data ecosystem simpler to use and safer to trust. This is a prerequisite for scaling AI responsibly in the years ahead.

**Deliverable:** Data Lake MVP (Month 19-20)

## 8. Reengineer High-Volume Workflows

Some of the City's most visible and resource-intensive services, for example, permitting, licensing, inspections, and scheduling, are built on processes that have developed over time, often becoming complex, manual, and inconsistent. These workflows are prime candidates for improvement, but not all are ready for automation as they stand. Applying AI to an inefficient or poorly designed process doesn't solve the problem; it amplifies the inefficiency, embedding bad habits into code and scaling the very issues AI was meant to fix, and may introduce added risk.

That is why Phase 2 emphasizes process redesign before automation. The City will use the insights and tools developed in early pilots to map, analyze, and simplify workflows before layering in AI. The City's working philosophy will be *redesign for automation*. This ensures that when automation or predictive tools are introduced, they enhance quality, speed, and consistency rather than replicating outdated procedures or inconsistent data handling at increased risk and increased cost.

Where workflows are ready, AI can deliver tangible improvements:

- **Automating document classification and routing** so staff spend less time sorting and more time resolving.
- **Prioritizing requests dynamically** based on urgency, impact, or regulatory deadlines.
- **Flagging anomalies or incomplete applications automatically**, reducing rework and delays.
- **Providing real-time visibility** into workloads and bottlenecks, helping managers allocate resources more effectively.

Each reengineered process will be treated as a model for replication. Considered as a practical example of how AI and human expertise combine to deliver faster, more predictable, more effectively, and with more transparent service. The first of these AI-optimized workflows should be live by Month 20, offering visible proof that responsible automation can benefit both residents and staff.

Over time, this disciplined approach of "**fix then automate**" will protect the City from one of the most common public-sector AI pitfalls: mistaking technology for transformation. By investing in process clarity and ownership before embedding AI, the City will ensure that automation reduces risk and builds resilience rather than unintentionally magnifying inefficiency or error.

**Deliverable:** First "AI-Optimized Process" (Month 20).

## 9. Implement Knowledge Management Systems

The City's ability to serve residents depends heavily on the experience of its people. Much of that knowledge lives in emails, shared drives, or the minds of long-tenured employees who understand the "why" behind each process. As retirements, role changes, and staff turnover continue, this institutional memory becomes harder to access, and the risk of losing it grows. Phase 2 will address that challenge by turning scattered experience into a shared and sustainable asset.

AI-powered knowledge management tools can help capture information that would otherwise disappear or be difficult to collect. The goal is to make the City's collective expertise easier to find, understand, and keep current. Activities include:

- Collecting and centralizing procedural documents, project lessons, and key communications into a searchable knowledge hub.
- Using AI to automatically tag, summarize, and connect related information so that staff can quickly locate what they need.
- Allowing employees to ask natural-language questions, such as "What's the process for emergency permit approvals?", and receive clear, sourced answers.
- Creating dashboards or prompts that help departments identify areas where documentation is thin or outdated.

This initiative also support succession planning and leadership continuity. When senior staff transition out, their experience and context will no longer leave with them. Structured knowledge capture will make onboarding new leaders faster and more consistent, reducing dependence on institutional memory and minimizing disruption.

The result will be a City organization that remembers what it learns, shares what it knows, and continues to grow more capable no matter how its workforce evolves.

**Deliverable:** Initial AI Knowledge Management System(s) (Month 24).

**Parallel Workstreams**

- **Governance Refinement:** Governance Committee evolves into an active oversight board reviewing scaling proposals.
- **Data Privacy Enhancement:** Policies evolve alongside integration to handle multi-department data sharing.
- **Change Management Continuation:** Training and cultural reinforcement continue as adoption broadens.
- **Performance Measurement:** Establish outcome-based metrics linked to service quality, ROI, and resident satisfaction.

**Expected Outcomes**

By the end of Year 2:

- A functioning, secure AI infrastructure integrated into targeted IT systems.
- AI improves service responsiveness in key departments.
- Staff in high-impact roles hold AI literacy or certification credentials.
- Residents have visibility into how AI decisions are made and monitored.
- Data and governance frameworks are mature enough to support advanced analytics and generative applications.

# Phase 3: OPTIMIZATION (24–36 Months)

"Institutionalizing AI for Continuous Improvement and Innovation"

## 1. Integrate Structured and Unstructured Data

In this phase, the City begins linking data across departments to enable more effective decision-making. The goal is not to integrate every dataset or system, but to focus on the high-value areas where better data connections can make a clear difference in service delivery and operational planning. These efforts will build on the data lake and governance foundation established in Phase 2.

The first priority will be to **combine structured data**, such as 311 service requests, infrastructure maintenance records, and public safety data, with **unstructured information**, including inspection notes, reports, and communications. When these sources are connected, AI models can begin to identify meaningful patterns and emerging

trends, such as where maintenance issues are likely to escalate or which neighborhoods may need additional public safety resources.

Some potential use cases might include:

- Predicting and prioritizing infrastructure repairs based on service requests, weather conditions, and asset data.
- Analyzing historical inspection and complaint data to identify potential health or safety risks before they occur.
- Supporting budget and staffing decisions by anticipating workload peaks or service bottlenecks.

The City's approach will be incremental and intentional, starting small and expanding only where clear value is demonstrated. By Month 28, the first set of cross-agency data connections and predictive models will be operational, providing insights that guide resource allocation and policy planning.

**Deliverable:** High-Value Modeling Pilot and Cross-Agency Data Integration (Month 28)

## 2. Operationalize Knowledge Management

At this stage, the City's knowledge management tools have moved from pilot to everyday use. What began as an effort to capture scattered information will become an active part of how employees work, learn, and share expertise. The goal is to make organizational knowledge dynamic with it being consistently updated, easy to access, and directly embedded into daily workflows.

AI tools will help organize and surface information at the point of need. Employees will be able to find procedures, historical records, and reference materials quickly, reducing dependence on manual searches or individual memory. As documents are created or updated, the system will automatically capture key details, tag related topics, and link to prior cases or policies.

This approach strengthens both succession planning and onboarding. New hires will spend less time tracking down answers, while experienced staff can focus on solving problems instead of repeating explanations. Over time, this living knowledge base will ensure continuity even as staff change, creating a smarter, more resilient organization that learns collectively rather than starting over with each new generation of employees.

**Deliverable:** Operational Knowledge Management System Fully Embedded in Daily Workflows (Month 32)

## 3. Scale the First Transformative AI Project

With governance, data, and foundational systems in place, the City will be ready to expand from targeted automation toward system-wide innovation. This phase represents a shift from improving individual functions to creating connected, intelligent services that enhance how the City operates as a whole. The emphasis will be on scaling AI projects that deliver measurable, cross-departmental impact and directly improve the resident experience.

Example transformative efforts could involve:

- Predictive public safety and emergency response, using data from multiple sources to anticipate incidents, allocate resources, and improve response times.
- Adaptive traffic and energy optimization, applying AI to real-time data from sensors, utilities, and transportation systems to improve flow, reduce waste, and support sustainability goals.
- Personalized resident portals and proactive communication tools, allowing citizens to receive tailored updates, reminders, and service recommendations based on their needs and past interactions.
- Smart infrastructure maintenance, combining sensor data and AI forecasting to detect early signs of road, bridge, or utility system degradation before costly failures occur.
- AI-supported community health and social services, integrating data from public health, housing, and social programs to identify at-risk populations and coordinate outreach more effectively.

These examples illustrate what is possible. However, the City will need to identify its own logical use cases based on need, data readiness, risk, and benefit. Each selected project will be guided by clear goals such as faster response times, improved equity, reduced costs, or better resident engagement.

The first citywide AI-enabled transformation project will go live by Month 33, demonstrating the practical value of coordinated, ethical, and transparent AI innovation across multiple City functions.

**Deliverable:** Citywide AI Transformation Project (Month 33)

## 4. Expand Civic Partnerships and Innovation Ecosystem

To sustain momentum and foster creativity, the City will broaden its network of collaborators and become an active participant in the regional AI innovation ecosystem. Public-sector innovation thrives when governments, universities, startups,

and established technology providers work together to experiment and co-develop solutions that address shared challenges. During this phase, the City will seek partnerships that bring in new perspectives, technical expertise, and access to research and emerging tools.

Examples of engagement activities may include:

- **Collaborating with local universities** on applied research and student projects focused on civic data, responsible AI, and public service innovation.
- **Partnering with startups and technology firms** to pilot small-scale AI solutions that can later be scaled citywide if successful.
- **Hosting an annual "Civic AI Challenge" or "Open Data & Innovation Week"** to invite the community, developers, and entrepreneurs to use City data for public good projects.
- **Establishing innovation partnerships with peer cities or regional agencies** to share lessons, avoid duplication, and explore joint initiatives.

These partnerships will create a healthy cycle of experimentation, feedback, and co-learning, keeping the City's AI approach fresh and adaptive. The first Civic AI Challenge will be launched by Month 34, marking the City's commitment to open collaboration and community-driven innovation.

**Deliverable:** Civic AI Challenge, and Innovation Partnership Framework (Month 34)

## 5. Institutionalize AI Oversight and Continuous Improvement

By year three, the City's AI program will move beyond pilots and scaling to optimizing. The existing AI Governance Committee must evolve based on lessons learned regarding governance in years one and two.

The Oversight Board will establish a rhythm of regular ethics and bias reviews, examining how AI is used, where it is delivering value, and where adjustments are needed. Reviews will include both internal audits and external consultation, where appropriate, particularly for high-impact or resident-facing systems. Findings will be shared transparently through public reporting cycles that summarize progress, highlight lessons learned, and outline next steps for improvement.

This process turns oversight into a learning process, not just an enforcement activity. The City will publish its first Annual AI Accountability Report, capturing performance metrics, resident sentiment, risk management results, and compliance milestones. It should be delivered to mark the completion of the roadmap period, showing measurable progress toward responsible, equitable, and effective use of AI across City operations.

**Deliverable:** Annual AI Accountability Reports (Month 36)

## 6. Enhance AI Resilience and Compliance

As AI becomes part of the City's everyday operations, reliability and accountability must grow alongside innovation. This phase focuses on ensuring that AI systems are not only effective but also secure, recoverable, and compliant with evolving laws and standards. The goal is to embed resilience directly into its AI ecosystem so that critical tools and data can withstand system failures, cyber incidents, or regulatory changes without disrupting essential services.

Work during this stage will include developing AI-specific continuity and recovery procedures, ensuring that models, datasets, and algorithms can be restored quickly and safely when needed. Regular model validation and stress testing will confirm that systems remain accurate and unbiased as conditions evolve, helping identify issues before they impact operations or residents.

At the same time, the City will continue to align its practices with state and federal AI regulations, maintaining transparency, privacy, and ethical accountability in every deployment. By Month 36, the City will have established a comprehensive AI resilience and compliance framework that safeguards the integrity of its systems while preserving public trust.

**Deliverable:** AI Resilience, Compliance, and Continuity Framework (Month 36)

**Parallel Workstreams**

- **Workforce Evolution:** Expand certification programs; create AI practitioner tracks for analysts, managers, and developers.
- **Resident Engagement:** Launch continuous feedback loops, including surveys, listening sessions, and community dashboards.
- **Model Lifecycle Management:** Standardize monitoring, retraining, and decommissioning protocols.
- **Performance Metrics and Reporting:** Formalize KPIs for ROI, satisfaction, equity, and resilience.

**Expected Outcomes**

By the end of Year 3:

- Aurora operates a mature AI ecosystem embedded in city operations.

- Governance and ethics reviews are institutionalized and public.
- Cross-department data integration enables proactive and equitable decision-making.
- The City is recognized as an AI innovation hub for public sector transformation.
- Staff are confident, capable, and continuously learning thereby sustaining the City's AI future beyond initial investment.

## Conclusion

# Appendix

## Glossary

- Structured data – Data that fits neatly into a table or database that is easy for computers to search and analyze (e.g., names, addresses…)
- Unstructured data – Data that doesn't have a fixed format (e.g., photos, videos…)
- Data lake - A data lake is a centralized repository that stores large volumes of raw data in its native format — structured, semi-structured, or unstructured — until it is needed for analysis.
- Interoperability - the ability of different systems, devices, or software applications to work together **—** to exchange, understand, and use information effectively without special effort from the user.