

Responsible Artificial Intelligence (AI) Use Framework

April 2026

A guide to the responsible adoption of AI in your municipality.



This guide was prepared with the assistance of AI.

Table of Contents

Executive Summary	5
1.0 Background	8
1.1 Introduction	8
1.2 Municipal Context: Key Drivers for this Framework	8
1.3 Purpose of this Framework	9
1.4 Key Takeaways	10
2.0 What is AI and Why Do We need a Framework?	11
2.1 What is AI?	11
2.2 What AI can do for us	13
2.3 What AI is not good at	13
2.4 Where We See AI Emerging	14
2.5 What are the risks with AI?	15
2.6 Municipal Obligations	16
2.7 Countering the Concerns	18
2.8 Why Municipalities Need a Framework	19
2.9 Key Takeaways	20
3.0 The Foundations for Responsible AI Use	21
3.1 Accountability and Responsibility	21
3.2 AI Governance	22
3.3 Set Guiding Principles: Your Guardrails for Using AI Responsibly	22

3.4	The Red Line	28
3.5	Experiment, Learn and Adapt	28
3.6	Key Takeaways	33
4.0	Getting Started	34
4.1	Why You Should Start Now (Call to Action)	34
4.2	Laying the Path for Change.....	35
4.3	Adopt a Phased Approach to Adoption.....	37
4.4	Suggested Training by Phase	44
4.5	Tools and Technologies	45
4.6	Build Your AI Tool Registry	46
4.7	Measure and Communicate Value.....	48
4.8	Key Takeaways	50
	Appendix A: Common Use Cases.....	51
	Appendix B: Training Resources.....	54
	Appendix C: AI Use Case Decision Tree	58
	Appendix D: Procurement Checklist for Solutions with AI	59
	Appendix E: Municipal Prompt Library	62
	Appendix F: Draft AI Policy.....	69
	Appendix G: AI Incident Response	73
	Appendix H: Sample One-Page Incident Report.....	75
	Appendix I: AI Q&A.....	76

Executive Summary

Purpose

This framework provides municipalities with structured, actionable guidance for adopting Artificial Intelligence in a manner that is responsible, compliant, and strategically sound. It is designed to help municipal leadership move beyond ad hoc AI experimentation toward a governed, value-driven approach that protects public trust while capturing real operational benefits.

Why This Matters

AI is already in your organization. The question is whether it's being used responsibly. Survey data from over 200 municipal staff reveals that more than 75% are already using AI tools without formal guidance, policy, or oversight. This exposes organizations to significant risks around data privacy, regulatory compliance, and reputational harm.

At the same time, the regulatory environment is tightening. Recent updates to MFIPPA and the introduction of the Enhancing Digital Security and Trust Act (EDSTA) now require municipalities to implement formal AI governance and data controls. Non-compliance is no longer a theoretical risk; it carries legal and reputational consequences.

What Could Go Wrong Without Action

Municipal leaders must be clear-eyed about the risks of both inaction and unchecked adoption:

Privacy and data leakage: Staff using consumer AI accounts may inadvertently expose personal or confidential information to third-party model training pipelines.

Accuracy and accountability: AI can produce confident but incorrect outputs ("hallucinations"). Without human review, errors can enter the public record.

Algorithmic bias: AI systems trained on biased data can produce inequitable outcomes, particularly in service delivery, enforcement, and hiring.

Cybersecurity gaps: AI introduces new attack vectors, including prompt injection and model manipulation, that require proactive security controls.

Loss of public trust: Uncontrolled AI use in decision-making undermines the community confidence that municipalities depend on.

Critical Insights

The framework offers three foundational insights for senior leadership:

Waiting has a cost. Municipalities that delay engagement with AI will face a steeper learning curve and widening capability gaps relative to peers who adopt early. The cost of entry today is low; it will only increase.

Governance is the prerequisite, not the afterthought. A cross-functional AI Governance Team, spanning IT, legal, HR, privacy, and operations, is essential before broad adoption begins. Without it, risk management is reactive rather than preventive.

The returns are real and measurable. Think of it this way: if your team saves 15 minutes here and 30 minutes there, it adds up quickly. Modest AI adoption can return \$3–\$4 in value for every \$1 spent on licensing^{1,2} and that's before you factor in the harder-to-measure wins like better quality work and less frustration.

The Roadmap: Three Stages

The framework recommends a phased, risk-tiered adoption strategy structured around three stages:

Stage 1 – Explore: Establish governance, procure enterprise-grade AI licenses, identify 5–10 departmental champions, and define clear "red lines", absolute prohibitions on AI use involving sensitive or personal data.

Stage 2 – Norm: Validate 2–3 pilot use cases, expand to 20–30 users, measure outcomes, and refine policy based on real-world learning.

Stage 3 – Scale: Embed AI into standard workflows, broaden access across the organization, and build AI literacy as a core institutional capability.

All use cases must be classified by risk tier (Low, Medium, or High) with corresponding approval requirements, disclosure obligations, and oversight mechanisms.

¹ <https://chamber.ca/re-igniting-canadas-productivity-engine-with-ai/>

² <https://www.toronto.ca/legdocs/mmis/2021/gl/bgrd/backgroundfile-171710.pdf>

The Bottom Line

The question for municipal leadership is no longer whether to engage with AI, but how to do so responsibly. This framework provides the governance structure, guiding principles, phased roadmap, and practical tools to enable municipalities to capture meaningful operational value, while protecting residents, staff, and the public trust that underpins everything municipalities do.

The time to act is now. This framework provides the roadmap.

1.0 Background

1.1 Introduction

Never before has a new technology moved as fast, or been as broadly adopted as Artificial Intelligence has.

From the first emergence of ChatGPT 3.2 in late 2022, we have seen an arms race between the companies responsible for ChatGPT, Gemini, Claude, CoPilot and various generative AI products vying for market share. The state of the art is changing daily. Things are moving so fast it can be bewildering, confusing and scary. Yes, there is undoubtedly too much hype, an AI bubble even, but we believe that there may be practical and useful applications of this new technology for municipalities with limited resources.

While not typically an early adopter of cutting-edge technologies, municipalities and municipal staff can see the very real value of AI and are jumping on board. As you will see our research suggests that more than 75% of staff in municipalities surveyed are already using AI – without formal guidance or guidelines.

We wouldn't give a member of staff a chainsaw to use before briefing them on the risks and training them how to use the saw safely. We should ensure that we do at least that for AI.

With the ubiquitous availability of public AI tools, alongside the increasingly pervasive integration of AI into commonly used software, ignoring AI is not really a viable option; but being an early adopter also comes with clear risks. So responsible experimentation, learning and iteration is the recommended strategy that municipalities should be encouraged to take.

And that's where this framework comes in – it is intended to provide a primer, along with practical guidance, principles and processes that can help municipalities pursue a strategic and responsible approach to adopting AI.

1.2 Municipal Context: Key Drivers for this Framework

In the Fall of 2025, we surveyed a sample of municipal staff across over 100 municipal organizations to understand how AI was being used.

The results identified that:

- Over 75% of staff are using AI tools for a variety of tasks, in most cases without any organizational guidelines

- 7% suggested their organization had a policy or guidelines, or training
- 15% responded that they had received training.
- 60% of staff are using AI outside of work for personal activities – like vacation planning, creative projects, learning and as a replacement for search.
- At work, respondents indicated they are typically using AI to assist with writing, editing, summarization, research, content creation, and to access technical support.

While many staff are actively embracing new systems and technologies, the majority are doing so informally, without guidance or guidelines. This raises real concerns around data privacy, quality control, and consistency of how these tools are currently being used.

At the same time, there is increased regulatory controls and scrutiny being placed on municipalities by the Province. In response to the evolving landscape, the Municipal Freedom of Information and Protection of Privacy Act (MFIPPA) has been updated, and the Enhancing Digital Security and Trust Act (EDSTA) has been introduced to require municipalities to implement effective governance and controls for handling personal information, addressing cybersecurity, and AI use.

At the same time, municipalities face unique pressures and challenges that make adoption of AI attractive and necessary. The challenge of “doing more with less” is not just a cliché but a lived reality for municipal councils, leaders and staff who must balance rising service expectations and demands with their own responsibilities and capacity. AI offers an attractive pathway to enhancing capacity without additional personnel costs for resource-constrained organizations, offering ways to work smarter, automating receptive tasks, and leveraging technology to increase efficiency.

But municipalities are of course also in a unique position regarding public trust – they are under increased scrutiny around how decisions are made and their communities are demanding transparency around decision making. Embracing new tools like AI comes with significant risks of harm and diminishment of community trust if used without clear guardrails, oversight, or shared understanding of how to mitigate the risks.

1.3 Purpose of this Framework

This framework is designed to provide background information and guidance to municipalities on how to embrace AI in a responsible, consistent and manner that is compliant with applicable municipal privacy and AI regulations.

This document was prepared and published in February 2026. With things moving so fast in the AI space, the framework and this guidance is intended to be a living document. One that must continue to be updated and evolve as AI continues to advance, as we learn more, as legislation is updated, and as municipal needs and the marketplace shifts.

This framework will be reviewed and republished on a periodic basis, updated with emergent good practices and shared lessons to keep it relevant and useful over time.

1.4 Key Takeaways



We recommend that municipalities in Ontario should:

- **Be open minded:** Accept that AI will continue to evolve and be willing to learn and adapt
- **Learn what's possible:** Understand current capabilities and limitations, learn by doing through examples and pilots
- **Be thoughtful and intentional about your use of AI** (organizationally and individually): Pause before using AI, consider, consider potential impacts on people, data, and trust



We recommend that municipalities in Ontario should not:

- **Ignore AI and hope it goes away:** avoiding AI increases unmanaged risk, uncoordinated and uncoordinated use creates risk and inconsistency
- **Panic:** AI adoption does not require immediate transformation, but careful experimentation is recommended
- **Worry:** Change is inevitable in this world and job losses outside of certain sectors have not materialized.. Focus instead on learning, reskilling and augmentation

2.0 What is AI and Why Do We need a Framework?

2.1 What is AI?

While there are many definitions of Artificial Intelligence (AI), for the purposes of this framework we are referring to systems that can perform tasks that normally require human intelligence, such as understanding context and meaning, learning from examples, recognizing patterns, and making predictions or recommendations.

2.1.1 Generative AI / Large Language Models

Generative AI is a subset of AI that use Large Language Models (LLM) to create (*generate*) new content – for example text, images, software code, or audio – based on patterns learned from large datasets. Common examples of these tools include ChatGPT, Microsoft Copilot, Claude, Google Gemini, and GovAI.

LLM's are good at many things, and we are still learning how to use them well. But typical uses include:


- Drafting emails, reports, briefing notes
- Summarizing long documents into plain language
- Rewriting content for clarity, accessibility, or for a certain audience
- Extracting key points from notes
- Brainstorming ideas
- Reviewing data and identifying patterns

AI tools also have capabilities to handle of voice to text, image recognition, and OCR.

2.1.2 Agentic AI

Agentic AI is an emerging set of AI capabilities that can run independently to design, execute, and optimize workflows with little human intervention or completely autonomously³. While LLMs respond to a prompt and typically stop, Agentic AI can act, orchestrate, coordinate, and continue until a task is complete.

³ <https://www.pega.com/agentic-ai?>



Agentic AI can support:

- Automated reporting and dashboards
- Monitoring announcements and summarizing changes
- Preparing background research
- Internal workflow coordination
- Monitoring and alerting
- Acting on triggers (e.g. classifying a record/document)

As such, Agentic AI can be thought of as a sophisticated set of automation tools with artificial intelligence built in.

2.1.3 Traditional AI

Machine Learning (ML), Predictive Analytics and Computer Vision are all examples of more traditional AI capabilities that have been around for decades. Though not as flashy as more recent breakthroughs, algorithmic ML is what powers many features in your phone, for instance – face recognition in your photo app is likely powered by ML. Though older, and perhaps in use in various existing business processes, algorithms and algorithmic decision making should follow the same guiding principles as we discuss

throughout this document.

2.1.4 Consumer vs. Enterprise Accounts

Not all AI accounts carry the same risk, and understanding the difference between consumer and enterprise AI accounts is important for determining how to responsibly use AI in the municipal context.

Consumer / free AI accounts (like the free versions of ChatGPT, Claude, or Gemini) are designed for personal use. These free accounts reserve the right to use whatever you input into them for model training and improvement. This means your conversations and data may be used to make the AI better. These accounts don't include contractual data protections, admin controls, or guarantees about how your information is handled.

Enterprise / Business AI accounts are paid versions of the same AI platforms that come with contractual agreements. These agreements guarantee that your data will not be used for training the model, and they include data retention controls (how long the data is kept) and admin oversight (organizational control). You get additional security and compliance protections built in.

As a rule of thumb: the more sensitive the information (such as in the case of municipal use), the more important it is to understand the data handling terms of whatever AI tool you're using.

2.2 What AI can do for us

For most, Generative AI tools are best understood as a “co-pilot”, a helper, an assistant that helps management and staff work more efficiently, prepare and communicate better, process information more effectively, identify new questions and perspectives, generate ideas and explore options.

AI is great at processing vast amounts of text, it's good at brainstorming, and its brilliant at pattern recognition – it can see patterns that are not immediately obvious.

2.3 What AI is not good at

AI is not an autonomous decision-maker and must not replace professional judgment or accountability. While AI tools are very good at generating human-sounding language, we must remember that they are not human and don't actually know or understand the subject matter in the text they produce; they are generating information based on text patterns that they have been trained upon.

So, despite how it may seem and however confidently AI asserts information, AI outputs cannot be guaranteed for accuracy, completeness, neutrality, or legal compliance. AI outputs should always be verified and be treated as suggestions, not answers. Human judgment and accountability remain essential.

2.4 Where We See AI Emerging

Obviously, the rise of Generative AI has been primarily seen via the emergence of Chatbots such as CoPilot, Gemini, ChatGPT, Claude.

But this is not the only place that we are seeing AI. AI is being embedded into all the M365 apps and platform, into Android and iOS, Google Workspace, and Windows and Mac operating systems.

In addition to these areas, we see the emergence of generalized products such as GovAI are built on top of the existing products features and controls, and we see the emergence of specialized LLM's tailored to specific use cases – such as reviewing the building code (e.g. Trex). We also see the use of AI in the reviewing of development and permit drawings for compliance purposes.

We are increasingly seeing AI being embedded into Municipal solutions. For instance, your recruitment solution may offer to summarize and rank applicants, your planning solution may summarize a planning application into a couple of sentences, your Web Content Management system might offer to improve the content you want to publish to the web.

We see dedicated solutions that incorporate image recognition and processing, such as video-based tools that scan the roadway for potholes or that OCR and route invoices for processing.

We see many municipalities (e.g. King, Markham) using customer facing Chatbots to provide enhanced customer support and service out of hours.

We are starting to see Agentic AI – that use the computer on your behalf – orchestrating activities across systems and processes – and these tools are being offered in large platforms – like M365, but we expect to see those filter down.

The key message here is that the IT industry is sprinkling AI into everything; municipalities should be aware of that, and be ready to ask critical questions, and better understand how AI is being used, what controls and ownership that you have and what it is doing with your data.

Real Examples

Calgary – Uses AI screening tools to streamline recruitment with human oversight for final decisions.

Winnipeg – AI-powered invoice automation.

Edmonton – Tracks wildlife using AI for public safety and environmental management.

Guelph – AI for ALPR parking enforcement.

London – AI tool for eligibility decisioning for homelessness prevention.

Halton Police – SARA system handles non-emergency calls.

2.5 What are the risks with AI?

While AI offers many benefits, municipalities must also be aware, acknowledge and actively work to mitigate the associated risks with using AI. These include:

- **Accuracy and “hallucinations”** – GenAI can produce confident but incorrect or incomplete answers. All outputs must be reviewed and verified by staff before being used in decisions, public communication, or formal records.
- **Privacy and confidentiality** – Pasting sensitive content into consumer/free GenAI accounts can expose confidential or personal information, particularly where it relates to residents, employees, enforcement activities, or negotiations.
- **Security and data residency** – External tools may store prompts and outputs outside Canada or in ways that conflict with municipal obligations. Any integration with core systems (M365, Finance, HR, Laserfiche, CityWorks, etc.) must go through appropriate security and privacy assessment.
- **Algorithmic bias and fairness** – Because GenAI relies on large, sometimes biased training data, its outputs may reflect stereotypes or skewed patterns. Use must be monitored for fairness, especially in hiring, enforcement, and access to services.
- **Records, transparency and accountability** – Outputs that influence decisions may need to be captured as part of the official record. Municipalities cannot rely on opaque, “black-box” decision-making; staff must remain accountable and able to explain how recommendations were developed.
- **Over-reliance on AI**: If staff become overly reliant on AI without reviewing the output, quality may decline and errors may get published. Maintaining an organizational culture of “AI provides the draft, humans provide the decision”, regular training and supervision can help combat this.
- **Staff expectations and change fatigue** – Over-promising on “AI” can create anxiety about job security or lead to disappointment if early pilots are poorly scoped. Communication must emphasize that GenAI is a tool to support staff, not to replace them.

When AI Goes Wrong

Fabricated Citations –

Newfoundland and Labrador’s Education Accord’s 10-year roadmap included 15+ citations from non-existent journal articles, none were caught before publication.

Weaponized Misinformation – KICLEI’s AI chatbot targeted 14 Canadian municipalities with climate misinformation, successfully convincing Thorold to temporarily withdraw from climate programs and Lethbridge to cut emission targets in half.

Government Service Fail –

The CRA’s AI chatbot provided accurate tax answers only one-third of the time, according to the Auditor General.

All these risks point to the importance of having a human-in-the-loop to verify the veracity of the information generated by AI. There have been numerous incidents around the world in the press that have been reputation damaging to their respective organizations – and all are a result of not being suitably aware of these risks.

2.6 Municipal Obligations

Accountability and responsibility for compliance with these obligations remains with the organization and individuals in the municipality that use AI tools.

So, the following municipal obligations continue to apply to our use of AI:

2.6.1 Records Management, MFIPPA and FOI

Content created by AI may or may not be subject to municipal records controls – and thus to the public disclosure, retention, archiving and destruction policies that apply to other types of content.

It is important to note that not every AI interaction constitutes a municipal record. Exploratory prompts, brainstorming sessions, and draft ideas generated with AI can generally be treated as transient working materials — similar to handwritten notes or whiteboard sketches — and may be disposed of without formal retention.

However, when AI-generated content informs a decision, is incorporated into an official communication or report, or is released publicly, it should be treated as a municipal record and managed accordingly.

So, a good rule of thumb is if AI-generated content informed a decision, was shared externally, or became part of a work product that carries the municipality's name or authority — it's a record.

The emerging best practice guidance (including from ARMA) suggests:

- Store the **output**, not the prompt history, unless the prompt materially shapes the meaning of the output
- Where AI content is incorporated into a document, report, or communication, that final document is the record — the AI exchange itself may not need to be retained separately
- Where AI is used to make or support a decision (e.g. a bylaw, a procurement recommendation), retain both the output *and* enough context to explain how it was generated (the prompting that got you to the output)
- Do not rely on the AI platform's own chat history as your records system — export and store the interaction, in your municipal EDMS or recordkeeping system

Your municipal Records Manager / Clerk should be able to provide advice here.

2.6.2 Data Protection and Privacy

Firms that build and train AI models require a supply of new data to continue and refine that training. By default, any data entered into a generative AI tool is added to the training data for the underlying model. This means that the model consumes the text and it becomes part of the LLM. This is particularly true of “free” versions of the tools – the trade-off being that users can use the tools, but in doing so they are supplying training data to the companies.

So, without controls that prevent information from being used in this way, municipal staff may “leak” information however unintentionally to these models. Typically, additional controls are available in paid or enterprise versions of AI software, but not in personal accounts even if they are paid accounts.

So, the advice here is for municipalities to avoid staff needing to use consumer or enterprise-based AI solutions, by providing an approved alternative solution that should be used.

Municipalities should note and communicate to their staff the importance of the Red-Line discussed later in the document – reinforcing for staff that confidential, sensitive, restricted content and any PII should never be shared in a consumer AI solution.

Of course, all organizations have specific obligations when handling Personally Identifiable Information, and this is especially true when using AI. Other commercially or politically sensitive information has the potential to leak if municipalities are not crystal clear about how their staff should use AI when handling this kind of information.

2.6.3 Fairness and Inclusion

Because Large Language Models have been trained on human generated content – literature, internet content, etc. – it has unavoidably picked up human biases when it comes to race, gender, nationality, sexuality, community, and viewpoints which can in turn affect decision making when those outputs are used.

In many cases these biases are subtle and difficult to discern. While the AI firms have worked to counter overt biases, the underlying models can and do exhibit biases. Individuals must be aware of that and make efforts to eliminate or reduce those biases in solutions under the municipality’s controls.

2.6.4 Cybersecurity

Alongside Privacy, municipal staff should also be aware and proactive about addressing cybersecurity concerns that come along with AI. In contrast to traditional software, AI has additional vulnerabilities such as:

- **Model Manipulation:** Where outputs are influenced to produce harmful or incorrect results.
- **Prompt Injection:** Where malicious inputs manipulate the system behaviour.

Municipalities where AI systems engage with public users and external vendors are particularly vulnerable to these threats.

Prior to implementing any new AI systems, these risks can be managed through a Privacy Impact Assessment (PIA) and a security review to ensure that the use case aligns with municipal cybersecurity standards.

For more advanced or high-risk application, an AI Maturity Assessment may also be needed to confirm the municipality has adequate governance, capacity, and staff readiness to manage the tool. There should also be a data retention and deletion plan for each AI system which outlines how long data will be kept, where it will be stored, and how it will be securely disposed of when no longer required.

2.7 Countering the Concerns

As we mentioned in the introduction, the emergence of AI has also been accompanied by a broad set of fears about the role of technology in society; not least the growth and dominance of a handful of major technology firms, loss of privacy and an emergent surveillance culture, concerning and cavalier business practices, and the ever-increasing control they wield.

But, related to AI specifically, municipal staff are understandably concerned about:

- **Job Losses** – AI is expected to automate many tasks faster than workers can retrain or transition. This raises concerns about unemployment, wage pressure, and widening inequality.
- **Loss of Human Control** – AI systems can make decisions that are difficult to explain or audit. When errors occur, it is often unclear who is responsible or how to correct the system.
- **Bias and Discrimination** – AI can reinforce existing biases embedded in historical data. This can lead to unfair or discriminatory outcomes at scale, especially in hiring, finance, and public services.

- **Privacy and Surveillance** – AI relies on large volumes of personal data, often collected without meaningful consent. This enables persistent surveillance and weakens individual privacy protections.
- **Misinformation and Trust Erosion** – AI makes it cheap and easy to generate convincing fake text, images, and video. This undermines trust in media, institutions, and even firsthand evidence.
- **Safety and Alignment** – AI systems can behave in unexpected ways when deployed in real-world conditions. At scale, even small misalignments between goals and outcomes can cause significant harm.

Organizations need to actively address and manage the fears and concerns that Councils, leaders, staff and their communities hold around the use and application of AI by committing to working transparently, communicating intent, putting in place the controls and guardrails to ensure that AI is used appropriately.

2.8 Why Municipalities Need a Framework

All the benefits, risks and fears discussed in this section point to why municipalities need a framework to ensure that:

- We stay compliant with legislation and obligations
- We learn how to use and apply AI in a carefully considered manner
- We help teams adopt and adapt to new tools and capabilities
- Human control is always maintained that qualified humans and not algorithms are making important decisions that affect people and communities
- Ensure that AI systems behave in ways that we understand/expect through verification
- Biases are prevented from impacting decision making
- We treat personal and other data in our care with the attention and due diligence our communities expect

2.9 Key Takeaways



We recommend that municipalities in Ontario **should**:

- Take a responsible approach
- Be aware of the risks and concerns
- Put in place guardrails to protect your organization, your teams and your staff



You **should not**:

- Delegate accountability or decision-making authority to AI systems
- Use AI in ways that undermine transparency, fairness, or explainability
- Deploy AI in public-facing or decision-making contexts without full consideration, review, and approval

3.0 The Foundations for Responsible AI Use

We have laid out the background on what AI is, what the benefits are, and we have identified the risks and the fears. How then should municipalities go about responsibly adopting AI?

We believe that this comes down to:

- Setting clear accountabilities and responsibilities,
- Establishing suitable AI governance,
- Defining clear principles to guide AI use,
- Defining a clear red-line – things that must not be done with AI, and
- Adopting an iterative test and learn approach.

3.1 Accountability and Responsibility

First, we should be clear that municipal councils, and senior leaders are accountable for what their organizations and staff do. Councils should ensure that the administration has done their due diligence and has an AI policy and guidelines in place.

The Administration, CAO and Senior Management Team, have overall accountability for how the organization uses AI – approving frameworks, policies and higher risk AI use cases, based on recommendations. This does not mean that the CAO / SMT need to become *the* experts in AI – but they are accountable for ensuring that the municipality has the right controls and governance in place, and they should ensure that they have established the framework to ensure that staff know what they can and cannot do.

Whoever leads IT / technology should be responsible for ensuring that any technologies in use at the municipality (including AI) comply with IT and cybersecurity requirements. They can only do that with the cooperation of their partners in business areas bringing forward new opportunities and solutions for review and with the support of the Senior Management Team, which is essential.

In addition to this responsibility, each member of staff, supervisor, and manager has a responsibility to be open and curious about AI and how it could be applied to benefit the municipality, but everyone does so alongside a personal responsibility to make themselves aware and compliant with AI policies, guidelines and best practices.

3.2 AI Governance

Municipalities should establish an AI Governance team, that can actively manage the AI framework, policy, guidelines and operational practices that guide AI use in your organization. We have provided a suggested approach as part of this document, but of course governance can be modified for each organization's unique needs, resources and structures.

If these roles exist in your organization, we recommend that people that are responsible for IT, legal, HR, privacy, finance, procurement, alongside representatives from leadership and business areas be part of this governance team.

Once the team has worked through the initial setup – as described in section 4 below, a monthly 1-hour meeting should suffice.

Once established, all decisions related to AI, including those listed below, can be worked through at this group (or by delegating to individuals), and recommended, where necessary to CAO/SMT for approval.

- Review and update AI policy and guidelines
- Review and approval of new use cases (following the AI decision tree)
- Review and approval of new tools that incorporate AI
- Maintain a list/register of AI tools and AI use cases
- Updating policy and guidelines based on learning

3.3 Set Guiding Principles: Your Guardrails for Using AI Responsibly

Each organization should set their AI guiding principles – to provide direction on how the organization should use and apply AI. The guiding principles should form the foundation for AI use and assessing all AI use cases.

3.3.1 Suggested Principles

The principles presented here are based on best practices and shared values held across municipalities and also draws on recently published guidance from the Information and Privacy Commissioner of Ontario. They can be adopted as-is or adjusted to suit local needs as necessary.

The suggested principles are:

- **Accountability:** Staff always remain responsible for all AI outputs and decisions.
- **Human-in-the-Loop Always:** AI assists staff, but staff make decisions.
- **Fairness and Freedom from Bias:** AI must demonstrably treat all residents and staff fairly and equitably.
- **Privacy and Confidentiality:** When using AI, we must always protect personal and sensitive information.
- **Security:** When using AI, we must protect municipal systems and data from AI-related risks.
- **Transparency and Disclosure:** We must be open about when and how AI is used.
- **Reliability and Trustworthiness:** We only use AI where outputs can be verified and trusted.
- **Safety:** AI use must not create risks to public or staff safety.
- **Value-Driven:** AI must solve real problems and deliver measurable benefits.

We have provided additional detail about the meaning and implications of each of these principles below.

Accountability

Staff remain responsible for all AI outputs and decisions.

AI is a tool, not a decision-maker. The person using AI is accountable for verifying accuracy, ensuring quality, and taking responsibility for outcomes. If AI makes a mistake, the staff member who used it owns that mistake. This means:

- Never blame AI for errors - you are responsible
- Review and verify all AI outputs before use
- Be prepared to explain how AI was used in any decision
- Maintain professional judgment and expertise
- Document AI use when required for decisions affecting residents

In practice: If AI generates a report with incorrect data and you publish it without checking, you are accountable - not the AI. Adequate time and resources need to be provided to staff to be able to effectively be that human-in-the-loop. When considering any AI solutions, projected gains in productivity need to be balanced with the needed investment in quality assurance.

Human-in-the-Loop:

AI assists staff: staff make decisions.

AI should augment human capability, not replace human judgment, expertise, or accountability. Every AI output must be reviewed by a qualified person before it is used, published, or acted upon. This means:

- No automated decisions without human review and approval
- Staff with appropriate expertise must validate AI recommendations
- Adequate time and resources must be provided to staff to be able to effectively be that human-in-the-loop (to ensure that quality is not sacrificed to quantity)
- Professional judgment always overrides AI suggestions
- AI cannot replace human interaction with residents
- Staff can and should override AI when necessary

In practice: AI can draft a planning application summary, but a qualified planner must review it, verify accuracy, and approve it before it gets posted to the website.

Fairness and Freedom from Bias

AI must treat all residents and staff fairly and equitably.

AI systems can inherit biases from their training data, potentially leading to unfair outcomes. Municipalities must actively work to prevent bias and ensure fair treatment. This means:

- Test AI outputs for bias before deploying to residents
- Monitor AI recommendations for patterns of unfair treatment
- Ensure AI works for diverse populations (languages, abilities, literacy levels)
- Provide non-AI alternatives for those who need them
- Be especially vigilant when AI affects access to services or individual rights

In practice: Before using AI to screen permit applications, test it with representative examples to ensure it doesn't unfairly disadvantage certain applicant types, neighborhoods, or demographics.

Privacy and Confidentiality

Protect personal and sensitive information at all times.

AI use must never compromise the privacy of residents, staff, or confidential municipal information. This means:

- NEVER enter personal information (names, addresses, SINS, health records) into consumer accounts of AI tools
- NEVER enter confidential municipal information (legal opinions, negotiations, personnel matters) into consumer accounts of AI tools
- Only use approved AI tools with data protection guarantees
- Conduct Privacy Impact Assessments (PIAs) before deploying AI that touches personal data
- Ensure AI vendors cannot use municipal data to train their models

In practice: If you need AI to help draft a response to a resident complaint, remove all identifying information first, or use only approved enterprise tools with privacy protections.

Security

Protect municipal systems and data from AI-related risks.

AI introduces new cybersecurity risks that must be managed proactively. This means:

- Only use approved AI tools
- Consider limiting the use of consumer AI tools on municipal devices or networks
- Be aware of prompt injection and manipulation risks
- Ensure AI tools meet municipal cybersecurity standards
- Report any security concerns immediately

In practice: A resident email asking you to "use AI to process this attachment" - don't. This could be a prompt injection attack. Use only approved AI tools and be skeptical of external instructions about AI use.

Transparency and Disclosure

Be open about when and how AI is used.

Residents and staff have a right to know when AI is being used, especially when it affects them. This means:

- Disclose AI use in public-facing content (reports, communications, chatbots)
- Be honest about AI capabilities and limitations
- Don't hide AI use or pass off AI work as entirely human-created

- Make it easy for residents to reach a human if they prefer
- Use standard disclosure statements

In practice: If AI helped draft social media content, include: "This content was created with AI assistance and reviewed by municipal staff."

Reliability and Trustworthiness

Only use AI where outputs can be verified and trusted.

AI is not always accurate - it can "hallucinate" facts, miss important details, or provide outdated information. This means:

- Verify all AI outputs against authoritative sources
- Cross-check facts, citations, and data
- Use multiple sources to confirm AI recommendations
- Be especially careful with technical, legal, or financial content
- When in doubt, consult human experts

In practice: AI says a bylaw was passed in 2023 - verify this in your records management system before including it in a report. AI can be confidently wrong.

Safety

AI use must not create risks to public or staff safety.

Some AI applications, if implemented incorrectly, could create safety risks. This means:

- Never use AI for life-safety decisions without extensive validation
- Emergency communications must be verified by qualified staff
- Infrastructure/asset management AI must be validated by engineers
- Public-facing AI (chatbots) must have clear escalation paths to humans
- High-risk use cases (Tier 3) require CAO approval

In practice: AI can help draft an emergency alert about a water main break, but qualified emergency management staff must verify accuracy and appropriateness before sending.

Value-Driven

AI must solve real problems and deliver measurable benefits.

Don't adopt AI just because it's trendy. Every AI initiative should have a clear purpose and demonstrate value. This means:

- Start with the problem, not the technology
- Define success criteria before implementing AI
- Measure actual time savings, quality improvements, or cost reductions
- Stop using AI if it's not delivering value
- Prioritize use cases that benefit residents and staff, not just "look innovative"

In practice: Before purchasing an AI-powered planning tool, ask: What problem does this solve? How much time will it save? Is it worth the cost? Will it actually make things better for planners and applicants?

3.4 The Red Line

Alongside the principles there should be a clear message to all staff about what municipal staff absolutely cannot do. We've called this the red line.



The “Red Line”: A Non-Negotiable Boundary

Your teams should know that:

- Under no circumstances should confidential or personal information be entered into consumer accounts or unapproved AI systems.
- Under no circumstances should AI fully author final reports without human input and review
- Under no circumstances should AI replace human judgement, ethics, or decision making
- Under no circumstances should new AI tools, or tools that have AI in them be purchased without AI Governance review

As discussed earlier, the “Red Line” is in place because consumer accounts of AI tools are different than enterprise accounts. Once the data is input into a consumer account of an AI system, it leaves your control and can be stored in servers anywhere in the world. This data may also be used to train AI models, which could leak in unexpected ways.

3.5 Experiment, Learn and Adapt

Unlike learning how to use software for issuing permits or processing tax payments – where there is a clear set of steps that you can follow to achieve your task, working with AI is more nuanced and abstract. Staff need to learn by doing - trying, testing and experimenting with the tools to learn where and how it can be useful to them.

With the guiding principles and red-line clearly defined, municipalities should be in a position to encourage responsible experimentation with AI, and staff should know that leaders support responsible experimentation – as long as the guiding principles and the red-line is respected.

3.5.1 Risk-Based Approach

We suggest that municipalities start their experimentation and learning in low-risk areas.

AI use cases can and should be categorized into three tiers, analogous to a traffic light system. Low risk (green light) applications are typically internal uses with minimal consequences that involve only the individual staff member or their working group. Low risk activities do not require review or authorization if using an approved tool, but the outputs must still be reviewed for accuracy.

Medium risk (yellow light) activities will likely require review and approval from the AI Governance team, may need a disclosure of AI use, and of course outputs must be checked.

Finally, High risk (red light) activities may significantly impact residents and public trust if handled incorrectly. These activities will require AI Governance, a PIA and possibly CAO approval depending on the application scope and impact. Transparency and thus disclosure is also of high importance for these use cases.

Remember, regardless of how the use case is categorized, **you are always accountable for the content generated or decisions driven by AI systems.**

Low risk			
Definition	Approval	Examples	Requirements
For internal use only with minimal consequences. Errors made for these tasks are easily identified and corrected before any external impact. No sensitive, confidential, or personal information involved.	None, if using approved tools	• Drafting routine internal emails • Taking personal notes or creating personal to-do lists • Generating meeting agendas for internal team meetings • Brainstorming ideas for new initiatives; developing ideas • Summarizing meeting notes	• Staff must review all outputs for accuracy before using any outputs (the human-in-the-loop concept) • No confidential, personal, or sensitive data is used (if confidential, personal data is used, it is not considered a low-risk use)

		• Creating initial outlines for internal documents or non-public facing reports • Formatting or restructuring existing text	• Use only approved enterprise AI tools • Output remains within individual or immediate team
Medium risk			
Definition	Approval	Examples	Requirements
Public-facing content or internal documents that inform decisions, but where AI is used only to assist with drafting, outlining, or summarizing portions of the work. Final content requires human review, editing, and professional judgment. May affect workflows or public perception but does not directly impact individual rights or involve automated decision-making.	AI Governance Team and Department head review	• Drafting public facing reports • Drafting sections of public-facing communications (press releases, social media posts) with staff completing final review and approval • Summarizing background research to support a staff report, with analysis and recommendations written by qualified staff • Generating initial outlines for website content, with content specialists writing and reviewing final text • Creating first drafts of internal policy documents for expert review and substantial revision • Extracting themes from public consultation feedback for staff analysis	• AI must not generate complete, final documents—staff must substantially review, revise, and finalize all content • Disclosure should always be considered: "This content was created with AI assistance and reviewed by municipal staff" • Quality assurance by subject matter expert • Check for accuracy, bias, and alignment with municipal voice/standards

		Drafting procedural checklists for staff workflows (e.g., permit intake steps) subject to review	Document type and audience determine appropriateness (internal working draft vs. external/regulatory document)
High risk			
Definition	Approval	Examples	Requirements
Use that could significantly impact individual rights, access to services, or public safety. Involves processing personal or confidential information, legal or regulatory decisions, or any element of automated decision-making. Errors could cause material harm, legal liability, or erosion of public trust.	AI Governance Team and CAO approval	- AI-powered chatbots on municipal websites (directing residents to municipal services, programs, or information) - AI-assisted screening or ranking of permit, license, or rental housing applications - Automated eligibility assessments for social assistance, housing programs, or municipal services - AI-generated recommendations in Freedom of Information (FOI) responses, including suggested exemptions or redactions - AI analysis of enforcement data influencing compliance or penalties - Predictive systems for infrastructure failure or public safety risks (e.g., road conditions, water main breaks) - AI-powered public-facing tools that interact directly with residents on	- Privacy Impact Assessment (PIA) required - Initially CAO approval should be mandatory (may be delegated once processes are well established) - Human review and final decision-making by qualified professional with appropriate expertise - Ongoing monitoring for bias, fairness, and accuracy (e.g. quarterly review, periodic bias and model-drift spot checks)

		sensitive matters (e.g., bylaw complaints, service eligibility) Automated surveillance or monitoring systems (e.g., ALPR, video analytics) that identify individuals or behaviors	- Clear documentation of AI's role in decision process - Transparent disclosure to affected individuals where appropriate - Regular review and validation of AI system performance
--	--	---	--

While we have provided these levels for guidance, uses must always be looked at in context. It is wise to err on the side of caution. Think about the ramifications of errors or harms that could be caused by using AI.

If a use case spans categories, the higher risk tier applies. If a use case has direct impact on community safety, user/resident rights, financial sustainability or other material impact – it should default to a high risk. If a use case changes in scope (e.g. an internal tool is to be made available publicly), it should be re-evaluated.

Seeking approval for something that might be acceptable at a lower tier is preferable than moving forward without proper authorization and finding out later that higher-level approval was necessary. Municipalities should foster an environment where asking for advice regarding AI is encouraged.

See [Appendix A](#) for a list of common municipal Low, Medium and High Risk use cases.

3.6 Key Takeaways



We recommend that municipalities in Ontario **should**:

- Learn what's possible
- Set up an AI working group with representation from it, legal, procurement, business areas to make decisions
- Establish a thoughtful approach to Ai experimentation
- You should establish an Ai community that can share and learn together



We recommend that municipalities in Ontario **should not**:

- Assume AI is solely an IT responsibility
- Leave experimentation unmanaged or unsupported

4.0 Getting Started

4.1 Why You Should Start Now (Call to Action)

We believe that AI is a transformative technology – a paradigm shift. Like other technological paradigm shifts – the computer, email, the internet, the smartphone – we think AI will be hugely beneficial to municipalities.

There were of course concerns, risks and fears associated with the introduction of each of these technologies at the time – but municipalities experimented, learned, adapted and now in most cases could not operate without them. We suggest respectfully that municipalities should learn from that paradigm shift experience. The cost of entry today is curiosity and a bit of time. That cost will only go up. It will be harder to start tomorrow, and harder still the day after.

While there are risks and fears associated with AI, we believe that the benefits outweigh the risks. We think that organizations that engage with AI technologies will build an understanding of the capability of these tools, advancing into more sophisticated use cases that will drive real value to their organization.

Those who wait or are hesitant to adoption will miss opportunities and will be faced with a steeper learning curve as they play catch up with the municipalities who adopted AI early on. The gap between the two groups of adopters will widen and learning will get more complex as time goes on, especially given the pace at which the AI systems are evolving.

We know that learning something new is uncomfortable — a toddler falls far more often than they walk. But stopping after the first few falls guarantees they never learn. AI is no different. You will make mistakes. You will feel inefficient at first. That is part of the process. The advantage is that you can experiment privately and at your own pace. There is very little downside to trying, and an enormous downside to waiting.

Approach it with curiosity. Treat it as a skill to be developed, not a threat to be feared. Encourage your teams to experiment, share what they learn, and build confidence through practice. Those who engage early – even imperfectly – will be the ones who adapt fastest.

Our recommendation is to start now.

4.2 Laying the Path for Change

It is important to understand that introducing AI is not really a technology project - it's an organizational change that affects how people work, what skills they need, and how they see their roles. For municipalities with limited capacity, managing this change thoughtfully is essential to success.

As noted earlier, there are legitimate concerns about AI. Common staff concerns may include:

- **Job security:** "Will AI replace me?"
- **Skills gap:** "I don't know how to use this"
- **Additional workload:** "I'm already stretched thin", "I don't have time for this"
- **Trust:** "Can I trust what AI tells me?"
- **Relevance:** "This won't work for what I do"

In response to these concerns, we suggest that municipalities proactively manage the change by:

- **Being transparent, communicate early and often** – Communicate that AI is a tool to help staff work smarter, not to replace people. Be honest about what's changing and what's not.
- **Starting with willing volunteers / champions** – Begin with staff who are curious and supportive. Their success will influence skeptics more than edicts and mandates.
- **Showing, not telling** – Real examples from peers are more powerful than theory. Let champions demonstrate practical wins and how AI is helping them.
- **Acknowledging the learning curve** – It's okay to feel uncertain. Provide support, training, and safe spaces to experiment.
- **Celebrating wins** – Share stories of time saved, quality improved, or frustrations reduced. Make success visible

The goal isn't just to "roll out AI." It's to ensure that staff feel informed, capable, and supported through the transition, and use it when they have it. When done well, change management:

- **Reduces resistance and builds adoption:** People support what they help create.
- **Prepares you for bumps in the road:** Risk assessment helps you anticipate and plan for challenges.
- **Promotes ethical and responsible use:** Clear guidance prevents mistakes before they happen.
- **Creates sustainable change:** Not just a short-lived initiative that fades away.

This is about finding the right balance between human judgment and AI capabilities, not replacing one with the other.

You can't manage change if people don't know what's happening. So, a clear communication plan is your roadmap for keeping everyone informed and aligned throughout your AI journey. Think of it as answering the basics: Who needs to hear what? When? Through what channel? And who's responsible for making it happen? Don't overcomplicate it. A simple communication plan keeps your messaging organized, ensures nothing falls through the cracks, and makes sure the right people are delivering the right messages at the right time.

Below is a sample communication plan that your municipality can adapt to fit your context and timeline:

<i>Audience</i>	<i>Strategy</i>	<i>Key Message</i>	<i>Channel(s)</i>	<i>Date/Time</i>	<i>Who Owns It</i>
<i>All-Staff</i>	<i>Energized and engage the company with AI implementation.</i>	<i>We're embracing AI to work smarter, not replace people. We'll learn together.</i>	<i>All-hands, Monthly newsletter</i>	<i>Feb 2026</i>	<i>CAO + Comms</i>
<i>Council</i>	<i>Inform and seek support for AI framework adoption</i>	<i>AI offers real values but requires guardrails. Here's our responsible approach and ROI.</i>	<i>Council workshop, briefing note</i>	<i>Feb 2026</i>	<i>CAO</i>
<i>Senior Leadership</i>	<i>Align leadership on approach and expectations</i>	<i>You need to visibly support this, not just approve it. Here's your role in making it work.</i>	<i>SLT meetings, Follow-up one-on-one</i>	<i>Feb 2026</i>	<i>CAO</i>

Remember, change fatigue is real. Don't over-promise on "AI magic." Set realistic expectations, acknowledge the learning curve, and create safe spaces to experiment and even fail. The municipalities that succeed with AI are the ones that treat it as a journey, not a destination.

4.2.1 Importance of Leadership Buy-In

As a leader, please don't say, "I just don't get AI", "I am not interested in AI", "I don't have time for AI" or anything like that. We think that it's vitally important for CAO and SMT to:

- Support the careful adoption of AI tools for specific use cases that further the goals of staff and the organization.
- Engage in training sessions and self-education around AI tools including practical hands-on learning.
- Share their own learning journey
- Allocate budget and time for the initiative

You should also brief council on what you are doing, framing AI as serving the community better with existing resources, while being transparent about costs and expected benefits, addressing risks and concerns, and reporting back on progress, lessons learned and benefits achieved.

4.3 Adopt a Phased Approach to Adoption

Next the framework suggests a three-phase approach to AI adoption, patterned on successful strategies adopted by municipal leaders in the AI space. Each phase has clear objectives, activities, suggested timelines, and clear success criteria.

4.3.1 Overview: The Three Phases

Phase	Focus	Staff Involved	Investment Level
Explore	Awareness, set the guardrails, safe experimentation	Leadership + 5-10 champions	Low (licenses, time)
Norm	Validate use cases, build confidence	Champions + pilot teams (20-30 staff)	Medium (expanded licenses, training)

Scale	Embed into workflows, build capability	Broader organization (50-100+ staff)	Higher (full rollout, ongoing)
--------------	--	--------------------------------------	--------------------------------

The key idea here is that you can pause between phases, and you don't need to follow these timelines. You don't need to rush to Scale if Norm isn't working, and it's okay to extend Explore if your organization needs more time.

4.3.2 Phase 1: EXPLORE

This could occur over 2 -3 months, or substantially longer, depending on the organization.

Objective: Raise awareness, establish guardrails, identify champions, and create foundation for safe experimentation.

Who's Involved: CAO, Senior Management, IT Lead, AI Governance Team, 5-10 departmental champions

Suggested Activities

Leadership & Setup

- CAO reviews framework with Senior Management Team
- SMT discusses organizational priorities AI could support
- Establish AI Governance Team (monthly 1-hour meetings)
- CAO communicates to all staff: "Why we're exploring AI"
- Allocate budget for AI licenses/tools
- IT restricts consumer AI tools on municipal devices
- Procure and configure enterprise AI licenses (data protection ON, training rights OFF)

Build Capacity

- Recruit 1-2 AI champions per *willing* department (volunteers preferred) – we acknowledge that there may be those who are not interested, but we suggest that CAO's encourage all of their teams to get involved
- Deliver "AI Champions Workshop" for champions (3 hours, hands-on)
- Distribute AI framework document to all staff

- Create internal FAQ page/wiki with AI resources

Experiment Safely

- Champions test prompts using suggested prompts as a starting point which can be found in the prompt library (see Appendix D).
- Weekly check-ins with champions (15-30 min)
- Document 3-5 "quick win" examples with time savings
- Collect feedback and share champion stories
- AI Governance Team reviews for any incidents or concerns

Success Criteria for this phase

Before moving to Norm, you should have:

- 5-10 active champions using AI weekly
- 3-5 documented quick wins with measurable time savings
- Zero data leakage incidents
- AI Governance Team meeting monthly
- Staff awareness of AI initiative (not resistance)
- Leadership actively supportive (not just passively approving)

But, if the following conditions are happening, it's likely that you are not ready to push forward:

- Champions aren't using tools (need better staff training/support)
- Multiple data privacy incidents (need better guardrails/communication)
- Staff backlash or fear (need better change management)
- Leadership not engaged (need executive sponsorship reset)

Decision Point: CAO and AI Governance Team assess: Ready for Norm phase, or extend Explore?

4.3.3 Phase 2: NORM

This could occur over 3 - 6 months, or substantially longer, depending on the organization.

Objective: Expand user numbers, validate use cases, build organizational confidence, refine governance, and prepare for broader rollout.

Who's Involved: Champions, pilot project teams (20-30 staff), AI Governance Team, Department Heads

Key Consideration: Internal productivity uses (Tier 1) can proceed with department head awareness. External impact uses (Tier 2-3) require formal review, disclosure, and quality assurance. See Appendix A for examples.

Suggested Activities

Share Learning & Identify and Select Pilots

- Host "AI Showcase" event for all staff (champions present real examples)
- Conduct staff survey: "What AI use cases interest you?"
- AI Governance Team selects 2-3 pilot AI projects based on:
 - Clear operational pain point
 - Enthusiastic department sponsor
 - Measurable success criteria
 - Mix of Low and Medium risk use cases

Expand Use and Run Pilots

- Deliver "AI 101" lunch-and-learn for staff being added to original AI users' group
- Provide targeted training for pilot participants
- Participants use AI for defined use case over 4-8 weeks
- Track usage, outputs, time savings, quality metrics
- Weekly check-ins with pilot teams
- Document challenges and lessons learned
- AI Governance Team monitors for incidents

Evaluate & Decide

- Collect quantitative data (time saved, cost avoided)
- Collect qualitative feedback (surveys, interviews)
- AI Governance Team assesses each pilot:
 - Did it solve the problem?
 - Was it worth the investment?
 - Can it scale to more users?
- Document findings
- Decide: scale successful pilots, refine marginal ones, stop unsuccessful ones
- Update prompt library and use case guidance
- Brief CAO/SMT on results

Success Criteria for this phase

Before moving to Scale, you should have:

- At least 1-2 pilots showing clear value (ROI > 2:1)
- 20-30 staff using AI regularly with confidence
- Documented use cases, prompts, and best practices
- Refined governance process (decision tree working well)
- No significant incidents or privacy breaches
- Department Heads supportive of broader rollout
- Staff feedback generally positive (trending positive, not necessarily universal)

But, if the following conditions are happening, it's likely that you are not ready to push forward:

- Pilots show marginal or no value (reassess use cases)
- Quality concerns with AI outputs (need better training/verification)
- Staff resistance growing (need better change management)
- Governance bottlenecks slowing everything down (simplify process)

Decision Point: CAO and AI Governance Team assess: Ready for Scale phase, or extend Norm?

4.3.4 Phase 3: SCALE

This could occur over 6 - 12 months, or substantially longer, depending on the organization.

Objective: Embed AI into workflows, expand to wider organization, build sustainable capability, and make AI literacy a core organizational skill.

Who's Involved: All staff (phased rollout), AI Governance Team, Department Heads, HR (for training)

Suggested Activities

Expand Access

- Identify next wave of users (prioritize departments with successful pilots)
- Procure additional approved licenses based on demand
- Develop department-specific guidance and prompt libraries
- Roll out licenses to additional staff in cohorts of 20-30
- Deliver standardized training (2 hours: tools, Red Lines, use cases)
- Provide department-specific follow-up sessions (1 hour)
- Establish "AI Office Hours" (weekly drop-in support)

Embed & Sustain

- Integrate AI guidance into onboarding for new staff
- Update job descriptions to include "AI literacy" as desired skill
- Incorporate AI use into performance conversations (where relevant)
- Establish Community of Practice (monthly meetups for users)
- Create advanced training tracks for power users
- Document and share cross-departmental use cases
- Update intranet with AI resources, FAQs, success stories

Ongoing Governance

- AI Governance Team continues monthly meetings
- Review new use case requests and vendor proposals
- Monitor usage, incidents, and value metrics
- Update policy and guidelines as needed
- Share learnings with peer municipalities
- Conduct annual Privacy Impact Assessments for high-risk tools
- Complete annual security reviews of AI systems
- Assess ROI and budget for next year
- Review and update framework annually
- Survey staff on AI literacy and satisfaction

Success Criteria for this phase

Within this phase you should have:

- 30-50% of staff using AI tools monthly (varies by role)
- Documented time savings of 5-10% for regular users
- AI integrated into standard workflows (not "special")
- Self-sustaining community of practice
- Minimal governance bottlenecks (process is smooth)
- Staff see AI as helpful tool, not threat
- Council and public aware of transparent AI use

But, if the following conditions are happening, it's likely that you need to revisit your program and closely examine what challenges you are facing:

- Growing staff fear or resentment
- Champions burning out from lack of support
- Usage not growing despite expanded access
- Quality issues becoming more common (not less)
- Leadership support is words, not actions

- ROI declining or flat despite increased investment

4.4 Suggested Training by Phase

Effective learning doesn't happen through a single training session or announcement. Furthermore, learning how to use AI is more intuitive and sense based than rote learning of step 1, step 2, step 3 and so on. So, AI training and learning requires a thoughtful progression that moves people from awareness to capability to confidence.

Below is a sample of training activities that align with the three-phase approach to help your staff understand the framework, develop hands-on skills, and ultimately make responsible AI decisions as part of their everyday work.

Phase	Training Activity	Who	Time	Outcome
Explore	Kickoff orientation	All staff	30 min	• Awareness of framework & approved tools • The Red Lines (never do this) • How to use the Decision Tree • Who to ask for help
	Champions workshop	5-10 champions	3 hours	• Hands-on practice + coaching skills • Role-specific job aids (e.g., AI for budget variance analysis, planning applications, job postings, etc.)
Explore	Weekly check-ins	Champions	30 min	• Learn by doing with support • Peer coaching • Share what is and isn't working
	AI-Coaching	All staff	Variable	•
Norm	Pilot onboarding	20-30 pilot users	1 hour	• Just-in-time learning for use case

Phase	Training Activity	Who	Time	Outcome
	Office hours	Any user	Drop-in	On-demand problem solving
Scale	Cohort onboarding	New user waves	1 hour	Standardized rollout
	New hire orientation	All new staff	30 min	Embedded in onboarding
	Lunch & learns	Interested staff	Monthly	- Ongoing skill building - Lessons learned
	Framework update	Core rollout team	Annually	- Share successes/improvements - Incorporate new practices

The success of training depends less on perfect design and more on the commitment to regular touchpoints, accessible support, and ongoing opportunities to learn from experience. Start with the Explore phase, listen to what your staff needs, and adjust as you go. For more training resources, see Appendix B.

4.5 Tools and Technologies

You will need to decide what tools as an organization you will use. At the base level you will want a general AI Assistant that you can provide to staff. Here there are 2 options:

Option	Description	Pro's	Con's
1	A Single tool for everyone (e.g., Microsoft Copilot, Gemini)	Simpler governance, training, support	One-size-fits-all may not meet all needs

2	Hybrid approach (e.g., GovAI for most, Copilot for M365 power users, ChatGPT Enterprise for specialized teams)	Right tool for right job	More complex to manage, train, support

In making your decision, you should think about:

- What do you already pay for? (a subscription to M365 or Google Workspace may mean you already have Copilot or Gemini rights)
- What fits your budget?
- What handles Canadian data residency well?
- What's easiest to support with small IT team?

Our suggestion is to start with ONE tool for Explore and Norm phases. Add others only if clear need emerges in Scale phase. Your governance group can evaluate and select the recommended option.

4.6 Build Your AI Tool Registry

You need to maintain a simple list of ALL the tools that use AI in the municipality. A simple register as illustrated below will suffice.

Tool Name	Vendor	What It Does	AI Features	Users	PIA Done?	Risk Level
Microsoft Copilot	Microsoft	General AI assistant	Text generation, summarization	15	Yes	Low

Website Chatbot	Chatbot Vendor	Resident Q&A	24/7 answers	Public	Yes	Medium
CloudPermit (planning)	CloudPermit	Permit management	Summarizes applications	4 in Planning and Permitting team	Pending	Medium

This is important, because you can't govern what you don't track, it will help with budget planning, and ensure you avoid duplication; but most importantly it is a requirement of the Enhancing Digital Security and Trust Act (EDSTA).

An AI enabled tool should not be used unless it is recorded in the AI Tool Registry, and department heads are accountable for ensuring that they follow the processes identified in this framework and for registering any AI uses in the AI registry.

4.6.1 Evaluation Process for New Tools

The introduction of any new tools that incorporate AI should be subject to review by the AI Governance Team. We have provided a suggested AI decision tree in Appendix C, and this process should be managed by the AI Governance Team.

So, when a department/team wants to add an AI-powered tool, the AI Governance Team should:

- Use the AI Decision Tree (Appendix C)
- Review the Procurement Checklist (we have provided a suggested checklist in Appendix C)
- IT conducts a security review, and a Privacy Impact Assessment (Clerks, Legal, or Third-Party) is completed if necessary
- AI Governance Team approves
- Add to registry

4.7 Measure and Communicate Value

We invest in technology to gain benefits – better, faster, more efficient are the typical things that we most often look for. So, you spend \$1 to get back \$4 in value – this is the return on investment.

It is important to understand that not all value can be expressed as direct, cashable savings, most often value is returned in time savings for staff – 15 minutes here, 30 minutes there – which adds up and frees staff up to work on other things that need to be done.

Because AI costs can be high, it is important that municipalities track their investment and measure the value of their returns – the benefits that they achieve – so that they can clearly and transparently justify their investment. This is considered best practice.

4.7.1 What to Measure

We suggest that you keep tracking simple and automated where possible, but try to measure the following for your AI initiatives:

- Time saved per user per month (estimated)
- Number of active users
- Frequency of use
- Cost per user
- Cost avoidance (calculated time savings × hourly rate)
- User satisfaction (quarterly pulse survey)
- Quality improvements (fewer errors, better outputs)
- Confidence with AI tools (increasing over time)
- Use case diversity (are teams finding creative applications?)

Real Savings



Halton Hills – 6 CoPilot pilot projects saved an estimated \$175,000 worth of staff time.

4.7.2 Value Tracking Template

The AI Governance Team should maintain a simple spreadsheet to track the value of using AI at the municipality to capture the return on investment and to be able to communicate the value to Council, SMT and the community.

Use Case	Users	Monthly Uses	Time Saved (min/use)	Monthly Time Saved (hrs)	Value (\$50/hr)	Continue?
Meeting summaries	8	32	30	16	\$800	Yes
Draft social posts	2	40	15	10	\$500	Yes
Budget variance reports	3	12	45	9	\$450	Yes
TOTAL	13	84		35 hrs	\$1,750/month	
Annual Value: \$21,000 vs. Annual Cost (15 licenses × \$30/mo): \$5,400 = ROI: 3.88:1						

In the sample above, the organization is achieving savings in the value of \$21,000 per year against a spend of \$5,400 making an almost four-fold return on investment.

It is worth noting that not all use cases will work, AI is not the right solution for everything, outputs may be poor, and other unforeseen issues may arise. Don't be afraid of stopping or pivoting if a use case doesn't work well.

We suggest that you actively monitor and report on the value perhaps quarterly to SMT, annually to Council (as part of a broader technology update), and that you continue to regularly share success stories internally with your teams.

4.8 Key Takeaways



We recommend that municipalities in Ontario **should**:

- Get started with realistic timelines/expectations
- Prioritize low-risk, internal productivity use cases first
- Start with real operational pain points identified by staff
- Invest in change management, communication and training
- Measure value and communicate results
- Advance to higher-risk or public-facing uses only with governance approval



We recommend that municipalities in Ontario **should not**:

- Wait for perfect clarity or certainty before taking any action
- Deploy AI reactively in response to vendor pressure or public hype
- Treat AI as a standalone initiative disconnected from service priorities
- Over-engineer governance before learning through use
- Ignore human factors (fear, resistance, capacity constraints)
- Force adoption faster than organization can absorb

Appendix A: Common Use Cases

The following are examples of common uses of AI in municipalities that we have seen in 2025 / 2026, they are provided to give you inspiration and ideas about how and where you might start exploring AI yourself.

Use Case	Risk Level	Risk Assessment
Prepare meeting notes and actions, prepare for meetings	Low	Internal use only. No decisions made. Staff review before use.
Wargaming - preparing for meetings / anticipating reactions	Low	Internal planning. No external impact.
Web content governance: audit pages for readability, broken logic, duplication, and update suggestions; create consistent FAQs	Low	Staff review all suggestions before making changes.
Council reports: convert technical work into executive summaries, key messages, and "what Council needs to decide" sections	Medium	Experts review outputs. Council can ask for more detail.
Solid waste & fleet: optimize routing/scheduling, anticipate missed pickups, and support maintenance planning	Medium	Affects service quality, not safety.
HR support: draft job postings, interview questions, onboarding checklists, and training materials	Medium	Could introduce bias. Needs HR and equity review.
Records & information management: auto-classify documents/emails into records categories and suggest retention labels	Medium	Wrong categories could mean early deletion of records. Needs validation.
Procurement: scan RFP/RFI drafts for gaps; generate evaluation criteria; summarize vendor bids and contract terms for staff review	Medium	Could miss important terms. Needs procurement and legal review.

Use Case	Risk Level	Risk Assessment
Contract management: flag risky clauses (indemnities, data residency, AI terms), generate negotiation questions, and create plain-language summaries	Medium	Might miss risky language. Legal review required.
Policy and bylaw drafting support: generate first drafts, consolidate stakeholder comments, and produce "options + pros/cons" memos for review	Medium	Bylaws have legal force. Could miss requirements. Goes through full Council approval.
Translation & plain-language: translate and simplify public communications	Medium	Errors could confuse residents. Risky for legal notices or emergencies. Needs qualified translator.
Community engagement: analyze large volumes of survey/comments to detect themes and sentiment, and produce traceable summaries for staff review	Medium	Could misrepresent input or miss minority voices. Needs equity review.
Permitting/planning: summarize submissions, extract key attributes, identify missing information, and draft applicant follow-up requests	Medium	Could miss critical details. Affects property rights. Planning staff must review.
Inspections: generate inspection checklists and standard notes; summarize findings; propose follow-up actions for inspector validation	Medium	Could miss safety issues. Qualified inspectors must validate.
Service request intake: AI chatbots to answer common questions 24/7, reduce call volume, and route requests (web + SMS)	Medium	Public-facing. Could frustrate residents or misroute urgent requests. Needs accessibility testing and clear paths to staff.
AI to classify service requests (and sometimes analyze text/photos) to speed routing and response prioritization	Medium	Wrong classification delays urgent issues. Photo analysis raises privacy concerns. Needs monitoring for bias.

Use Case	Risk Level	Risk Assessment
Council/committee communications: auto-generate plain-language "what happened" briefs and podcasts from meetings/minutes for residents	Medium	Could misrepresent decisions. Affects public trust. Needs review before publication.
Contact centre support: speech/voice recognition and assisted responses for low-priority calls	Medium	May not work for all residents (accents, languages, disabilities). Could exclude vulnerable people. Needs testing.
Infrastructure/asset management: predictive maintenance and condition forecasting using historical work orders and sensor/IoT data	High	Wrong predictions could lead to failures (water mains, roads). Public safety risk. Needs engineering validation.
Transportation: traffic and transit optimization (signal timing, incident detection, ridership/pattern analysis)	High	Safety risk if systems fail. Could affect emergency response. Must ensure fair service across neighborhoods.
FOI / access to information: draft responses, locate likely relevant records, propose exemptions/redactions for human review	High	Legal risk under MFIPPA. Wrong exemptions could expose personal info or wrongly withhold public records. Needs expert review.
Emergency management: summarize incident logs, generate situation reports, draft public updates, and translate advisories quickly	High	Time-critical. Errors endanger public safety. Wrong info during emergencies risks lives. Needs immediate verification.

Appendix B: Training Resources

To help guide responsible use, the Treasury Board of Canada Secretariat developed the [FASTER](#) framework that can be adapted by your organization.

F	Fair	Ensure that content from these tools does not include or amplify biases and that it complies with human rights, accessibility, and procedural and substantive fairness obligations; engage with affected stakeholders before deployment.
A	Accountable	Take responsibility for the content generated by these tools and the impacts of their use. This includes making sure generated content is accurate, legal, ethical, and compliant with the terms of use; establish monitoring and oversight mechanisms.
S	Secure	Ensure that the infrastructure and tools are appropriate for the security classification of the information and that privacy and personal information are protected; assess and manage cyber security risks and robustness when deploying a system.
T	Transparent	Identify content that has been produced using generative AI; notify users that they are interacting with an AI tool; provide information on institutional policies, appropriate use, training data and the model when deploying these tools; document decisions and be able to provide explanations if tools are used to support decision-making.
E	Educated	Learn about the strengths, limitations and responsible use of the tools; learn how to create effective prompts and to identify potential weaknesses in the outputs.
R	Relevant	Make sure the use of generative AI tools supports user and organizational needs and contributes to better outcomes for clients; consider the environmental impacts when choosing to use a tool; identify appropriate tools for the task; AI tools aren't the best choice in every situation.

This modified framework from the Government of Canada, [Guide on the use of generative artificial intelligence](#), look at 8 concerns that are of particular interest to a public sector organization.

Concern	Description	Questions to ask
Impact on public trust	Generative AI can simulate human-like communication which can seem authoritative but may contain errors. If the results were to be used in decision-making, they could lead to unfair treatment of some people or to misguided policy. As well, when considering the use of these tools in public facing communications, it is critical that the government not share inaccurate information, which would contribute to misinformation and erode public trust.	• Could the use of AI cause a decrease in trust from the public? • Does the use of AI for decision making materially impact the public?
Protection of information	The protection of personal, classified, protected and proprietary information is critical when using generative AI systems. The suppliers of some generative AI tools may inspect input data or use it to further train their models, which could result in privacy and security breaches. Some tools, public or otherwise, may not meet the privacy and security requirements set out in legislation and policy.	• Could the use of AI expose private, personally identifiable, or proprietary information to the public or AI training data?
Legal risk	Risks could also arise from the opacity of generative AI models and their potential for producing inaccurate, biased or inconsistent outputs. This opacity makes it difficult to trace and understand how the AI system produces outputs, which can undermine procedural fairness in instances where a government institution is obliged to provide clients with reasons for administrative decisions, such as decisions to deny benefits. The quality of AI outputs can also impact individuals' legal rights. For example, biased outputs could lead to discrimination in services, potentially violating human rights.	• Could the use of AI put the County in a position of legal risk?
Transparency	Conversational agents or chatbots that use generative AI can produce responses that are so humanlike that it may be difficult to distinguish them from those of a real person. As a result, clients may think they are interacting with a human. Similarly, clients may think an email they have received was written by a person when it was actually generated by an AI tool. On the other hand, clients might think they are interacting with an AI tool when they are actually dealing with a real person. Transparency about whether a client is interacting with a person or a chatbot is essential to ensure that the client is not misled and to maintain trust in government.	• Does the use of AI need to be disclosed in the final product?

Accountability	Generative AI systems have capabilities that allow them to make assessments or determinations about clients as part of the service delivery process. For example, a generative AI system could be used to summarize a client's data or to determine whether they are eligible for a service. These administrative uses can affect how staff views and decides on a case, which has implications for the client's rights, interests and privileges. There needs to be an accountable human who takes responsibility for the outcomes of using AI for a particular purpose.	Is there an accountable human behind the use of AI for this purpose?
Bias	Generative AI tools can produce content that is discriminatory or not representative, or that includes biases or stereotypes (for example, biases relating to multiple and intersecting identity factors such as gender, race and ethnicity). Other sources that may contribute to biased content include data filtering, which can amplify the biases in the original training set, framing of the prompt, and model bias. Widespread use of these technologies could amplify or reinforce these biases and dominant viewpoints, and lead to less diversity in ideas, perspectives and language, as well as potential harms.	Could additional bias be introduced to a process or amplified by using AI?
Quality	Generative AI technologies can produce content that appears to be well developed, credible and reasonable but that is in fact inaccurate, nonsensical or inconsistent with source data. This content is sometimes referred to as a "hallucination." Also, content generated by AI tools may not provide a holistic view of an issue. Instead, it may focus on prevalent perspectives in the training data.	Could the quality of work be degraded by using AI?
Environment	Generative AI can contribute significantly to greenhouse gas emissions and water usage through the energy demands of model training, operation, and the production and cooling of supporting server infrastructure. While it has potential to aid in climate change mitigation, its benefits must be weighed against the urgent need to reduce environmental impacts.	What is the relative environmental impact of using AI for this purpose?

Each concern is given an impact rating which corresponds to a numerical value for calculation purposes. For each task, the user rates how much that concern worries them when thinking about implementing generative AI for

that task. The qualitative measures are Very High=10, High=5, Medium=3, Low=1. The impact rating and the concern rating can then be multiplied to generate a score.

Impact on public trust	Protection of information	Legal risk	Transparency	Accountability	Bias	Quality	Environment
Very High		High		Medium		Low	
10		5		3		1	


This results in a maximum score of 380 for an activity that is perceived as very high risk across the board, and a minimum score of 38 for an activity that is perceived as low risk across the board. An example of how an activity is scored might look like this:

Task: Meeting management

Concern	Ranking	Score	Rationale
Impact on public trust (10)	Low (1)	10	Tool used internally with human scrutiny. Final product edited and vetted by a human.
Protection of information (10)	Medium (3)	30	Confidential or private information may be shared during meetings and transcribed by tool.

Appendix C: AI Use Case Decision Tree

You can use the following decision tree to ask / answer questions about your AI use case to determine if it is a suitable use case to explore / implement.



Appendix D: Procurement Checklist for Solutions with AI

This suggested checklist is intended to be used by the AI Governance Team, or those delegated to assess new AI opportunities. You can use this checklist when:

- Buying any new software or service that includes AI functionality
- Renewing contracts for tools that have added AI features
- A vendor pitches an "AI-powered" solution
- Evaluating whether to enable AI features in existing tools

If a vendor mentions AI, machine learning, automation, or "intelligent" features, it's probably a good idea to review this checklist.

Data & Privacy Questions

Ask the vendor these questions and document their answers:

Data Collection & Use

- What municipal data will the AI access or process?
- Where is the data stored? (Must be in Canada or approved jurisdiction)
- How long is data retained?
- Can we request deletion of our data?
- **Red flag if:** Data stored outside Canada without clear legal agreement

Training Rights

- Will our municipal data be used to train AI models? The **answer must be NO**
- Can we opt out of data being used for model improvement?
- Does the vendor use our data to improve services for other clients?
- Are prompts/inputs stored separately from our municipal systems?
- **Red flag if:** Vendor can't guarantee our data won't train their models

Privacy & Confidentiality

- Vendor has completed our Privacy Impact Assessment (PIA) questionnaire
- Tool complies with MFIPPA requirements
- Personal information handling meets Ontario privacy standards
- Vendor will sign confidentiality agreement
- **Red flag if:** Vendor refuses PIA or can't explain MFIPPA compliance

Security & Access

- Data encrypted in transit and at rest
- Vendor has SOC 2 Type II or equivalent certification
- Multi-factor authentication available
- Role-based access controls included
- Vendor will complete our security questionnaire
- **Red flag if:** Vendor dismisses security questions or has no certifications

Municipal Obligations

- Tool supports AODA/accessibility requirements
- French language available (if required for your municipality)
- Complies with records management requirements
- Supports FOI request processes
- **Red flag if:** Accessibility is "coming soon" or French is not available

Vendor Accountability

- Vendor will provide AI incident response plan
- Clear escalation path if AI causes problems
- Vendor has liability insurance covering AI-related issues
- SLA includes response times for AI failures

- **Red flag if:** Vendor disclaims all liability for AI errors

Contract Terms (Legal Review Required)

- Contract explicitly prohibits using our data for training
- Data ownership clearly belongs to municipality
- Right to audit AI system and data handling
- Exit plan: how to extract our data if we cancel
- Indemnification if AI causes harm or privacy breach
- **Red flag if:** Contract has broad vendor disclaimers or unclear data rights

Appendix E: Municipal Prompt Library

The following are provided as inspirations for those thinking new to using GenAI tools.

Use Case	Prompt Example
Administration and Clerks	
Summarizing complex reports for executive briefings	Act as an executive assistant. Summarize the following [report/document] into a concise executive briefing. Include: • Executive summary (max 250 words) • Key findings and implications • Critical risks or opportunities • Recommended actions or decisions required [Insert or paste content here]
Drafting stakeholder communications	Draft a stakeholder communication regarding [topic]. Tone: [professional, urgent, educational...] Audience: [specify] Key message: [list 2-3 points] Length: [specify] Ensure the communication is clear, easy-to understand, accessible, and aligns with the Town's voice.
Documenting meetings	Convert these meeting notes into formal minutes. Include: 1. Attendees 2. Key discussion points 3. Decision made 4. Action items 5. Next meeting date [Paste notes]
Policy research and comparison	Act as a policy analyst. Research and compare how other Ontario municipalities handle [policy area such as short-term rentals, noise bylaws, climate action]. Focus on: 1. Different regulatory approaches and their rationales 2. Enforcement mechanisms and effectiveness 3. Lessons learned or challenges encountered 4. Recommendations for adapting best practices to our context

Use Case	Prompt Example
	Provide citations for specific municipalities referenced. [Review citations provided]
Council report presentation	Prepare a Council report on [topic] following standard municipal format: • Purpose: [state objective] • Background: [provide context] • Analysis: [key considerations, options, implications] • Financial Impact: [budget implications] • Recommended Motion: [draft motion text] Use professional municipal language and assume Council has general but not detailed knowledge of the subject.
Review a Council Report or presentation	Review the attached [report/presentation] and generate challenging questions that may be asked by [audience]. For each question: 1. Identify the specific concern or angle being raised 2. Suggest a clear, factual response 3. Flag any gaps in the material that should be addressed 4. Note any statements that could be misinterpreted or require clarification Provide 8-10 questions ranging from straightforward clarifications to more challenging critiques. [Paste content]
Communications and Public Affairs	
Crisis Communication Development	Develop a crisis communication framework for [situation - e.g., water main break, service disruption, emergency declaration]. Create: 1. Initial holding statement (for use within first hour) 2. Detailed public notice (for website and media) 3. Social media posts (3-4 updates as situation evolves) 4. FAQ addressing likely resident question. 5. Internal staff talking points Tone should be authoritative, calm, transparent, and solution focused.
Content Translation	Translate the following municipal notice into [language]. Ensure: 1. Plain language appropriate for general public 2. Cultural sensitivity and context

Use Case	Prompt Example
	3. Preservation of all key information and legal requirements 4. Natural phrasing rather than literal word-for-word translation Note: This is AI-assisted translation and should be reviewed by a fluent speaker before publication. [Paste content]
Social Media Content Calendar	Create a one-month social media content calendar for our municipality covering: 1. Key municipal events and deadlines 2. Seasonal themes (e.g., winter preparedness, summer programs) 3. Service reminders (e.g., waste collection, tax deadlines) 4. Community engagement opportunities For each post, provide: • Suggested copy (appropriate for Facebook/Twitter/Instagram) • Recommended posting day and time • Suggested image description or visual concept • Relevant hashtags [Paste content]
Plain Language Conversion	Rewrite the following technical or legal content in plain language suitable for general public understanding: • Target reading level: Grade 8 • Audience: Residents with no specialized knowledge • Maintain: All legally required information • Improve: Clarity, accessibility, approachability [Paste technical content]
Planning and Development	
Development Application Analysis	Analyze this development application against our Official Plan and Zoning By-law: • Application: [describe - e.g., 50-unit residential development on Smith Street] • Current designation: [specify] • Requested changes: [list]

Use Case	Prompt Example
	Provide: 1. Conformity assessment with Official Plan policies 2. Zoning compliance analysis 3. Key planning considerations (density, compatibility, servicing) 4. Potential concerns or issues requiring further analysis 5. Questions for the applicant or technical studies needed [Paste application details]
Public Consultation Summary	Synthesize this public consultation feedback into a structured summary: Organize by: 1. Major themes and how frequently each appeared 2. Areas of consensus vs. divergent opinions 3. Key concerns raised 4. Suggestions or recommendations from participants 5. Demographic patterns in feedback (if available) Maintain objectivity and represent diverse perspectives fairly. [Paste consultation responses]
Site Plan Conditions Drafting	Draft site plan approval conditions for [development type] addressing: 1. Landscaping and buffering requirements 2. Parking and access standards 3. Stormwater management 4. Building design and materials 5. Servicing and municipal requirements 6. Timing and securities Use standard municipal conditional approval language appropriate for legal agreements.
Finance and Budget Management	
Budget Variance Analysis	Analyze this budget variance data and provide insights focused on: 1. Key trends or patterns 2. Significant outliers or unexpected variations 3. Implications for year-end projections 4. Recommendations for corrective action or reallocation

Use Case	Prompt Example
	Present findings in executive summary format suitable for CAO or Council briefing. [Paste budget data]
Financial Policy Comparison	Compare our [policy type - e.g., reserve policy, debt management policy] against: 1. Provincial legislative requirements 2. Recommended practices from municipal finance officers 3. Policies from 3-5 comparable municipalities 4. Industry standards or benchmarks Identify: • Areas where our policy aligns with best practice • Gaps or areas for potential strengthening • Innovative approaches other municipalities have adopted • Recommendations for policy updates
Human Resources	
Job Description	Update this job description to reflect current role requirements: Original description: [paste] Updates needed: 1. Incorporate digital literacy and technology competencies 2. Emphasize collaboration and adaptability 3. Include equity, diversity, and inclusion responsibilities 4. Modernize language (remove dated terminology) 5. Ensure gender-neutral language throughout Maintain: Core qualifications and essential duties
Change Initiative Communication	Draft an employee communication about [organizational change - e.g., new software implementation, process change, restructuring]. Address: 1. What is changing and why 2. How it affects employees 3. Timeline and implementation approach 4. Support and training available

Use Case	Prompt Example
	5. Where to direct questions Tone: Transparent, reassuring, practical. Acknowledge potential concerns while emphasizing positive intent and support. Length: [specify]
Training Material	Create a training module on [specify topic] for [target audience] for a workshop that is [specify duration]. Include: 1. Learning objectives (3-5 specific, measurable outcomes) 2. Key concepts explained simply with relevant examples 3. Practical exercises or application activities 4. Knowledge check questions to assess understanding 5. Additional resources for further learning 6. Handout or job aid for reference after training Ensure accessibility and accommodation for diverse learning styles.
IT and Digital Services	
Technical Documentation for Non-Technical Audiences	Explain this technical concept or system to non-technical staff: Topic: [e.g., cloud migration, cybersecurity protocols, new software functionality] Create: 1. Plain-language explanation of what it is 2. Why it matters to daily work 3. What changes for users 4. Benefits and improvements 5. Common questions and answers Avoid jargon or define technical terms when necessary. Use analogies where helpful. [Paste technical details]
Knowledge Base Articles	Create a knowledge base article for this common technical issue: Issue: [describe problem] Affected systems: [specify] User impact: [explain]

Use Case	Prompt Example
	Article should include: 1. Problem description in user terms 2. Symptoms users might experience 3. Step-by-step troubleshooting instructions 4. When to escalate to IT support 5. Prevention tips if applicable Format for easy scanning with numbered steps and clear headings.

Techniques for Verifying and Validating AI Outputs

As we have noted, although AI tools can increase productivity, human judgement is still needed. Before any AI generated content is used or shared, it must be reviewed and verified.

There are a few suggested ways in which AI output can be verified:

Human-in-the-loop: This technique is the most important safeguard which involves a staff member actively evaluating, revising, and approving all AI outputs. In municipal settings, where errors may have serious repercussions, human review by specialists / technical experts is crucial. Accountability remains the obligation of the staff using the output.

SIFT Method: The SIFT method is a straightforward framework derived from information literacy best practices which suggests the following steps to test the veracity of information in media, that is also applicable to AI outputs; Stop, Investigate the Source, Find better coverage, and Trace claims, quotes and media to the original text⁴.

Cross-checking multiple AI Models: A practical technique for identifying errors or hallucinations is to cross-check the output of one AI system using another AI system. For example, if you put the same question in ChatGPT through Gemini and both tools produce similar answers, the confidence in the result increases. You can also check for hallucinations and fabricated sources through other AI models as well.

⁴ <https://guides.lib.uchicago.edu/c.php?g=1241077&p=9082322>

Appendix F: Draft AI Policy

The following provides a sample/draft AI policy which organizations will need to adjust as required to align with your policy standards and your organizational and operational realities:

AI POLICY

Approval date: [insert date]

Revision date: [insert date]

Review frequency: [insert review frequency – at minimum annually]

Purpose

This policy is the policy that documents the current approach to using Artificial Intelligence (AI) technologies within *[insert municipality]* (the “Municipality”). The purpose is to ensure that AI is used in a manner that aligns with organizational values, legal obligations, and societal norms.

Scope

This policy applies to all elected officials, employees, full-time, part-time and temporary staff, contractors, freelancers, and other agents who use or intend to use AI and LLM tools within the Municipality or using the data of the Municipality.

Policy statement

Technology resources are the property of the municipality and intended for the benefit of the municipality. The use of AI and LLM tools is subject to the specific policies and guidelines that follow

Principles

The following principles apply to the Municipality’s use of AI:

- **Accountability:** Staff always remain responsible for all AI outputs and decisions.
- **Human-in-the-Loop Always:** AI assists staff, but staff make decisions.
- **Fairness and Freedom from Bias:** AI must demonstrably treat all residents and staff fairly and equitably.
- **Privacy and Confidentiality:** When using AI we must always protect personal and sensitive information.
- **Security:** When using AI we must protect municipal systems and data from AI-related risks.

- **Transparency and Disclosure:** We must be open about when and how AI is used.
- **Reliability and Trustworthiness:** We only use AI where outputs can be verified and trusted.
- **Safety:** AI use must not create risks to public or staff safety.
- **Value-Driven:** AI must solve real problems and deliver measurable benefits.

Freedom of Information

The activity records and communications related to municipal IT systems and equipment is information that may have to be released to the public under the Municipal Freedom of Information and Protection of Privacy Act (MFIPPA).

Confidentiality

The sharing or access to confidential data must be limited to only that which is required for the vendor or third party to perform the work agreed to in the contract.

Tools

At this time, only the following AI tools are permitted. No other AI or LLM tools shall be used by anyone in the organization without prior authorization by the IT Governance Team / Clerk / CAO.

- [E.g. Microsoft CoPilot]
- [Insert approved tool]
- [Insert approved tool]

Governance

An AI Governance team will be established by the CAO. Any exception or deviation from the permitted tools shall require review and the written approval / permission of the AI Governance Team / Clerk / CAO. The AI Governance Team is responsible for defining processes for assessing and reviewing all AI use cases. The AI Governance IT Governance Team / Clerk / CAO will maintain a register of AI use cases. All other municipal policies must be followed when using AI tools, including security and privacy guidelines.

Regulatory Compliance

All provincial and federal regulatory compliance must be adhered to first, before considering the use of AI tools within the municipality.

The Province of Ontario has enacted the Enhancing Digital Security and Trust Act (EDSTA) which requires the Municipality to put in place governance and controls for AI use. The federal government has proposed the Artificial intelligence and Data Act (AIDA as part of Bill C-27).

All Ontario Law, including MFIPPA, FOI must be considered and adhered to when using AI tools.

This policy shall be reviewed annually or whenever relevant laws or regulatory compliance standards change

Disclosure

The Municipality will disclose publicly where AI has been used, with the following disclosure:

"This content was created with AI assistance and reviewed by municipal staff."

AI Incident Response

The Municipality will maintain an AI Incident Response process to ensure compliance with legislative requirements.

Exemptions

Any exemptions from this policy must be granted in writing and signed by the municipal CAO.

Noncompliance with this Policy

Appropriate action will be taken for failure to comply with this policy, which may include disciplinary action up to and including termination.

Definitions

For the purpose of this policy, the following definitions apply:

- **Artificial Intelligence:** Systems that simulate human intelligence processes such as learning, reasoning, and self-correction.
- **Confidential Information:** Information that is limited to only the audience it was intended for
- **Contractor:** Any third-party company or individual who is engaged by the Municipality on a temporary basis
- **Copyright:** The exclusive legal right, given to an originator or an assignee to print, publish, or share materials including software, and to authorize others to do the same
- **Information:** Data of any kind, raw or processed, confidential, sensitive or public
- **Information System:** A system that processes, stores or transmits data of any kind

- **IT Services:** Refers to employees and contractors who are dedicated to the installation, maintenance, operation and coordination of the municipality's IT infrastructure and services.
- **Large Language Model (LLM):** Advanced AI models trained on vast text datasets to understand, generate, and interpret human language
- **Machine Learning (ML):** A subset of AI involving algorithms that improve automatically through experience.
- **Municipal Freedom of Information and Protection of Privacy Act (MFIPPA):** Ontario legislation that provides the framework for municipal access to information and the protection of personal information.
- **Procurement:** The process of locating and agreeing to terms and purchasing goods, services, or other items from an external source
- **Sensitive Information:** Information of a personal nature, such as health information or performance evaluations.

Related Policies and Information

The following policies may have information related to this policy. Should you have any doubt as to which policy applies, contact Human Resources or the Clerk's office.

- Information Security Policy
- Password Policy
- Social Media Policy
- Federal and Provincial Regulations

Appendix G: AI Incident Response

You should have an incident response plan ahead of time that identifies how you will handle incidents. If/When Something Goes Wrong with AI:

Step 1: STOP (Right Now)

1. Stop using the AI tool
2. Take a screenshot of what happened
3. Tell your supervisor immediately

Step 2: Figure out How Bad Is It? (Within 1 Hour)

	Situation	Who to Contact	When
	• Personal info went into public AI tool • Wrong info published to residents • Discriminatory decision made	IT Lead + CAO + Clerk + Legal	Now
	• Wrong info published internally • Confidential info in enterprise tool • Biased output used in decision	IT Lead + Your Director	Within 4 hours
	• Minor error caught before use • Internal draft only	IT Director + Supervisor	Today

Step 3: Fix It (Within 24 Hours)

- **If personal info leaked:** Call the AI vendor to delete it. Legal assesses if we must report to Privacy Commissioner.
- **If wrong info was published:** Correct it immediately. Post explanation/correction/apology if needed.
- **If someone was treated unfairly:** Re-review the decision. Contact them with explanation and fix.

Step 4: Learn From It (Within 1 Month)

Fill out the one-page incident form (in Appendix H) and give to your organization's IT lead to review with the AI Governance team. The AI Governance Group will then:

- Figure out why it happened
- Update training or policy
- Share lessons learned

Give this form to your organization's IT lead within 1 week.

Remember:

1. Report fast – Don't hide mistakes, we fix and learn from them together
2. When in doubt – Contact your IT lead
3. No blame – Honest mistakes help us learn

For any questions, contact your IT leader.

Appendix H: Sample One-Page Incident Report

Your Name: _____ Date: _____

What happened? (2-3 sentences) _____

What AI tool were you using? _____

What information was involved? (Check one)

- Personal info (names, addresses, SIN, etc.)
- Confidential municipal info
- Public information only

Who was affected? (Check one)

- Just me
- My team only
- Residents/public → About how many? _____

What did you do?

- Stopped using the tool
- Told my supervisor
- Corrected published information
- Other: _____

Why do you think this happened? (Check all that apply)

- I didn't know the rule
- I made a mistake
- The tool didn't work right
- The policy wasn't clear
- Other: _____

Supervisor signature: _____ Date: _____

Appendix I: AI Q&A

Can I use AI tools without approval?

Only with **approved enterprise tools** on your municipality's authorized list. Never use free/public AI tools without IT Governance approval. Check with your IT lead for the current approved list.

What do I do if I accidentally input sensitive information into AI?

1. Stop using the tool immediately
2. Tell your supervisor and IT lead right away
3. Follow the AI Incident Response plan (Appendix G)

See Section 2.6 for obligations and Appendix G for detailed incident response.

Which AI tool should I use?

Check your municipality's approved tools list maintained by the AI Governance Team. Common options include Microsoft Copilot, GovAI, or other authorized enterprise tools. See Section 4.3 for tools and technologies.

A vendor wants to sell us an AI solution. What now?

1. Do not commit without proper review
2. Contact IT lead and procurement
3. Work with your IT lead to review the AI Vendor Management & Procurement Checklist (Appendix D)
4. Vendor must complete Privacy Impact Assessment
5. Requires AI Governance Team approval

See Appendix D for complete vendor evaluation process.

Can I use AI for [specific task]?

Use the AI Use Case Decision Tree (Appendix C) to determine:

Tier 1 (Green): Proceed with approved tools, review outputs

Tier 2 (Yellow): Department head review required

Tier 3 (Red): PIA and CAO approval required

Common use cases are listed in Appendix A.

How do I write good prompts?

See the Municipal Prompt Library (Appendix D) for tested examples. Key tips:

- Be specific about audience, length, and tone
- Provide context and examples
- Iterate and refine based on results

What if the AI output is wrong?

Always verify AI outputs before use - hallucinations are common. Cross-check facts, rely on your professional judgement, and never publish AI content without human review. See Section 2.3 for what AI is not good at.

Can I put resident information into AI?

No. Never input personal information (names, addresses, SINs, health records) or confidential municipal information into any AI tool. This is a "Red Line" violation. See Section 3.3 for complete Red Line policy.

Exception: Use of resident data in **approved, internally governed AI systems** may only occur where authorized and compliant with privacy, security, and legal requirements.

Do I need to disclose when I use AI?

- **Tier 1 (internal):** No disclosure required
- **Tier 2 (public-facing):** Yes, include: "This content was created with AI assistance and reviewed by municipal staff"
- **Tier 3 (high-risk):** Yes, with detailed documentation

See draft AI Policy in Appendix F.

Where do I get training?

Contact your HR team or department AI champion for:

- Framework review sessions
- Lunch and Learn sessions
- Hands-on tool training
- AI Office Hours

See Section 4.1 for getting started guidance.

Something went wrong with AI. Now what?

Follow the AI Incident Response plan (Appendix G). Report incidents promptly - we learn from mistakes together. The severity determines response:

- **Critical** (personal info leaked): Contact IT + CAO + Clerk + Legal immediately
- **Medium** (wrong info published): Contact IT + Director within 4 hours
- **Minor** (caught before use): Document and report today

Who do I contact for help?

- **Immediate issues:** Your IT lead
- **Policy questions:** AI Governance Team
- **Training:** HR team or your department's AI champion
- **Vendor questions:** Procurement team
- **Privacy concerns:** Clerk or privacy officer

Produced by



www.perrygroupconsulting.ca

Our reports are designed to meet accessibility guidelines, and we strive for accessibility and inclusivity.

~ ~ Trademarks acknowledged ~ ~