

TORONTO LEADS AGAIN

AI, Robotics & the Intelligent City

A human-governed framework for AI and robotics that expands capability while protecting rights, privacy, work, childhood and democratic authority.

Campaign Policy Proposal - Not a City of Toronto Publication



Jamie Atkinson for Mayor of Toronto

Public release | August 2026

People first. Families first. Machines serve. Democracy decides. Nature sets the constraint. Art gives the future meaning.

Contents

1. Why this matters
2. Executive Summary
3. Toronto Today and Inherited Work
4. Vision and Outcomes
5. Visualizing the Strategy
6. Strategic Priorities
7. Targets, Scenarios and Planning Assumptions
8. Implementation and Delivery
9. Accountability and Public Reporting
10. Sources and Notes

Why this matters

Teach a person to teach a robot

Give a person a fish and they eat for a day. Teach a person to fish and they eat for a lifetime. Teach a person to teach a robot to fish - while caring for the environment - and that person, their family and their community can prosper for generations.

The AI revolution should belong to Torontonians

Public policy should broaden access to capability while preserving human authority, rights, privacy, childhood, work and democratic accountability.

The signature idea

Give a person a fish, and they eat for a day. Teach a person to fish, and they eat for a lifetime. Teach a person to teach a robot to fish - while caring for the environment - and that person, their family and their community can prosper for generations.

The public principle

The AI revolution should not happen to Torontonians. It should belong to Torontonians.

Executive Summary

AI, Robotics & the Intelligent City is the governance framework for a machine-age Toronto. Its premise is that AI should expand human capability without replacing democratic authority, due process, privacy, professional judgement or childhood safeguards.

The signature opportunity is simple: the AI revolution should belong to Torontonians. Compute Commons, the Intellectual Lifeboat, Civic Intelligence Review, One Front Door, ethical proving sites and workforce pathways are

designed to broaden access to capability. Robotics is treated as physical capacity for housing, inspection, maintenance and dangerous work when it demonstrably improves outcomes.

The technology-governance rules in this report are paired with Report 17, AI Job Disruption, Worker Transition & Economic Security. That report focuses on employment effects, entry-level opportunity, worker navigation, paid transition pathways and responsible workforce change as AI adoption expands.

The first term establishes transparency, risk tiers, human accountability, procurement rules and bounded pilots before large-scale deployment.

Major commitments

- Publish a City AI and automated-system registry with responsible owners and purposes.
- Adopt risk tiers and impact assessment for consequential systems.
- Define prohibited or presumptively unacceptable uses, especially where rights and children are concerned.
- Create human escalation and accountable case ownership for resident-facing AI.
- Launch bounded ethical proving sites for robotics and AI with success and stop rules.
- Develop Compute Commons and Intellectual Lifeboat access pathways for residents, students, workers and small organizations.

How this fits the governing program

Financial treatment: first-term actions in this report must reconcile to Report 03. Unsigned partner funding and unverified technology savings are not booked in the municipal base case.

How to read commitments and numbers

Status	Meaning
OFFICIAL BASELINE	Current public evidence or an inherited program.
ADOPTED CITY TARGET	A target already approved by the City; not a new campaign promise.
CAMPAIGN COMMITMENT	A first-term action proposed by the campaign, subject to lawful Council/budget authority.
CAMPAIGN TARGET	A measurable result the campaign proposes to pursue.
PLANNING ASSUMPTION / SCENARIO	A decision-support input, not a guarantee.
PARTNER-DEPENDENT OBJECTIVE	Requires another government, agency, utility, institution or private partner.
LONG-RANGE ASPIRATION	Part of the 25-year direction; future Councils retain democratic authority.

Key public baselines and planning anchors

Indicator	Value	Date / status	Use in this report
Jobs in Toronto	~1.624M	2025 Official survey baseline	Current economic scale.
Net new jobs	23,410	2025 Official survey result	One-year result; do not extrapolate mechanically.
Applied compute decade band	20-80 MW	Decade scale Campaign capacity scenario	Requires real connections and utilization evidence.
Regional applied compute horizon	150-400 MW	Longer horizon Partner-dependent scenario	Utility/institutional/private and senior-government scale.
City AI registry	Planned Q3 2026	2026 Existing City initiative	Build on inherited governance, do not claim start from zero.

Human capability in the machine age Give a person a fish, and they eat for a day. Teach a person to fish, and they eat for a lifetime. Teach a person to teach a robot to fish - while caring for the environment - and that person, their family and their community can prosper for generations. The AI revolution should not happen to Torontonians. It should belong to Torontonians. The purpose of AI and robotics in this platform is not to replace the human city with a machine city. It is to expand what people can learn, build, repair, create, discover and care for. Students, tradespeople, workers, creators, entrepreneurs and public servants should be able to command powerful tools while humans retain purpose, judgment, responsibility and democratic authority. That principle connects the Intellectual Lifeboat, Compute Commons, construction robotics, AI enablement, Civic Intelligence Reviews, the Economic Watering Well and the Ethical AI & Robotics Proving City. Technology is the multiplier. Human flourishing is the objective. Nature sets the constraint.

Toronto Today and Inherited Work

Current Conditions

This strategy begins from the programs, infrastructure and conditions Toronto already has. Existing work is identified as the inherited baseline; new campaign commitments are stated separately. Baseline facts that control the plan Baseline What the evidence says City AI governance The City's responsible GenAI guidance directs staff to use approved Microsoft Chat and requires verification; a first Algorithmic Impact Assessment is developed and an AI registry is planned for Q3 2026. [AI2][AI3] Ontario ecosystem Vector reports 121 new AI companies, 5,813 AI jobs created, 71,167 professionals requiring AI skills and 2,371 \$85k+ jobs for AI-related graduates in Ontario in Apr. 2025-Mar. 2026. These are Ontario-wide, not Toronto-specific. [AI1] Federal compute AI Compute Access Fund up to \$300M for SMEs; Sovereign Compute Infrastructure Program can provide about \$890M to infrastructure build layer in the current call. [AI4][AI5] Power constraint Toronto data centres were

~293 MW in Sep. 2023; new Ontario connection interest is orders of magnitude larger, making energy/public-value discipline necessary. [GRID6][GRID7] Workforce Ontario construction may face a 27,300 worker gap by 2035; robotics should be evaluated as augmentation plus new skilled work, not a magic labour replacement. [JOB2] Throughout this report, official baselines, campaign calculations, campaign proposals and partner-dependent decisions are identified separately.

Vision and Outcomes

Technology policy succeeds when human capability and democratic authority grow together.

- AI assists; people decide.

- Residents know when consequential automated systems are used and who is accountable.
- Public technology expands access to knowledge and capability.
- Workers are trained and involved before automation is scaled.
- Robotics improves safety and productivity where real-world evidence supports it.
- Children receive stronger privacy and developmental safeguards than ordinary commercial users.

Visualizing the Strategy

Selected campaign visuals from the earlier reference reports are retained here because they explain relationships that are harder to understand in prose alone. They are explanatory illustrations; the surrounding text and tables define commitments, costs, authority and status.



Ethical AI Proving City: test bounded use cases, publish results and scale only when rights, safety and public value are demonstrated.



Physical AI should augment trained workers, not bypass safety, accountability or professional judgement.



Human capability remains the point: AI should expand what creators, workers, students and public servants can accomplish.

Strategic Priorities

The technology strategy pairs broad access to useful AI with enforceable rules for risk, privacy, human accountability and physical robotics.

1. AI registry + risk tiers + Algorithmic Impact Assessment

Residents should be able to know when City government uses consequential AI and what safeguards apply. Toronto would maintain a public inventory of material AI systems, classify them by risk and require proportionate impact assessment before higher-risk uses proceed.

The framework would identify purpose, data sources, accountable human owner, affected rights, testing, monitoring and appeal or correction pathways. Prohibited or unacceptable uses would be stated plainly rather than left to procurement-by-procurement interpretation.

2. One Front Door / Civic OS with human ownership

Residents should not need to understand City organizational charts to obtain a service. One Front Door and the Civic Operating System are service-design concepts: connect requests, status information, institutional knowledge and accountable staff ownership across departments while keeping human authority clear.

AI may help route, summarize and retrieve information, but it does not become the decision-maker simply because it is efficient. Accessibility, multilingual service, records management and privacy are built into the service architecture.

3. Intellectual Lifeboat and Civic Intelligence Reviews

The Intellectual Lifeboat is the campaign's promise that access to useful intelligence will not be reserved for large companies and wealthy households. Libraries, schools, community institutions and workforce partners can help residents learn to use AI, understand its limits and find human help when technology changes work or services.

Civic Intelligence Reviews apply the same philosophy inside government: model options, surface evidence, compare accessibility and lifecycle effects, preserve institutional memory and make assumptions visible before major decisions are locked in.

4. Compute Commons for residents/students/creators/SMEs

AI opportunity increasingly depends on access to computing resources as well as software. Compute Commons would create bounded pathways for students, researchers, creators, startups and civic projects to reach shared compute without requiring every participant to buy expensive infrastructure.

Access would be governed by budgets, energy constraints, security, fair-use rules and public-value criteria. The program would start small and scale only if utilization and outcomes justify the cost.

5. Ethical Compute Exchange readiness with energy/security/provenance rules

Where multiple institutions need secure access to models, data or compute, Toronto can help convene common governance rather than create an unregulated marketplace. An Ethical Compute Exchange would define provenance, cybersecurity, privacy, energy reporting and contractual accountability for participating public-interest uses.

The concept is partner-dependent and would require universities, health institutions, governments and private operators to agree on roles. It is not a municipal data pool and would not override existing privacy or sectoral law.

6. Robotic housing/maintenance/inspection proving sites

Physical AI can expand construction, inspection and maintenance capacity when robots are designed around real work, skilled supervision and measurable safety. Toronto would use bounded proving sites to test tasks such as material handling, repetitive fastening, inspection or maintenance where automation can reduce injury and improve throughput.

Pilots would compare total installed cost, quality, worker safety, reliability and schedule performance against conventional methods. Workers and trades remain central to setup, supervision, exception handling and quality assurance.

7. Safe Passage bounded child-safety pilot and standards work

Children's safety deserves innovation but also unusually strong limits on surveillance. Safe Passage is a bounded standards-and-pilot concept for improving safety around schools, crossings and routes through environmental design, operations and carefully limited sensing where justified.

No pilot should normalize facial recognition or continuous child tracking. Privacy impact assessment, public notice, data minimization, independent evaluation and clear stop conditions are prerequisites.

8. Slipstream machine geography under study

Slipstream describes the machine geography of delivery: the physical and digital systems that allow approvals, utilities, logistics, inspections and construction to move in a coordinated sequence. It is intended to reduce avoidable waiting and conflict rather than weaken democratic or safety review.

Initial work would map bottlenecks in selected housing and infrastructure programs, test shared scheduling and identify where automation or common data can remove repetitive administrative work. Any broader system would be evidence-led and human-governed.

-

Targets, Scenarios and Planning Assumptions

Planning Assumptions and Key Calculations

Show the arithmetic and its limitations so optimistic assumptions cannot quietly become promises. Calculation register Calculation / model Method / interpretation Compute scale discipline A 120-140 kW rack is a fundamentally different infrastructure class from a workstation or robot edge computer. The report therefore presents compute in layers rather than aggregating all AI into a single MW target. Data-centre queue 10,000+ MW of proposed Ontario connections is ~40% of a recent ~25 GW Ontario peak. A queue is not a utilization forecast and should be screened for economic/strategic value.

Personal-robot stress case Illustrative only: 500,000 small devices at 0.5-1.5 kWh/day use 0.09-0.27 TWh/year (average 10-31 MW). If charged in a four-hour window, gross coincidence could be ~63-188 MW before diversity/smart charging. The lesson is charging/distribution standards, not a need for a dedicated robot power plant. Ontario AI job context 5,813 AI jobs / 71,167 professionals requiring AI skills = ~8% as many new AI jobs as the broader skilled-professional base in Vector's Ontario snapshot; most AI-era value therefore depends on AI capability spreading into existing occupations, not only creating narrow AI job titles. [AI1] Planning arithmetic is used as decision support. Where official forecasts or project-specific engineering supersede a campaign screening calculation, the official or engineered value governs.

Implementation and Delivery

First 100 Days

The opening 100 days establish the public AI inventory, risk rules, accountable ownership and a small number of bounded pilots before any large-scale deployment is contemplated.

- Publish the initial AI/automated-system inventory and named owners.
- Propose risk tiers, impact-assessment triggers and human-escalation requirements.
- Identify prohibited-use categories and interim procurement safeguards.
- Select a small number of bounded proving-city pilots with measurable public outcomes.
- Launch an Intellectual Lifeboat access and training design with libraries, education and workforce partners.

2027-2030 First-Term Action Plan

During 2027-2030, Toronto expands useful AI access and carefully evaluated robotics while enforcing risk tiers, privacy, cybersecurity, resident rights, workforce transition and public reporting.

Phase	Action	Lead / authority	Resident result
Register	Public AI and automated-system inventory	City / agencies	Residents can see purpose, owner and risk class.
Assess	Risk tiers and impact assessments	City / privacy / legal / program owners	Consequential systems face stronger review.
Protect	Human decision rights, privacy and child safeguards	City / partners	Automation cannot silently become final authority.
Prove	Bounded AI and robotics pilots	City / agencies / labour / vendors	Scale follows measured public value and safety.
Enable	Compute Commons and Intellectual	City / libraries / universities /	More people can use advanced

Phase	Action	Lead / authority	Resident result
	Lifeboat	partners	tools productively.

2031-2051 Long-Term Direction

Over the long term, Toronto should be a place where residents broadly share the benefits of AI and robotics while democratic authority, privacy, human authorship and the right to meaningful human review remain intact.

- Build a civic intelligence layer that augments institutional memory and analysis while preserving human authority.
- Expand physical AI only after safety, workforce, lifecycle and public-value evidence is strong.
- Keep standards adaptable as models, hardware, cybersecurity threats and law change.

Risks and dependencies

Risk	Specific mitigation
Privacy or civil-liberties harm	Apply data minimization, impact assessment, access controls, auditability and stop authority.
Vendor lock-in	Require portability, exit rights, open interfaces where practical and procurement competition.
Workforce displacement	Pair pilots with worker consultation, training and transition plans before scale.
Model error or bias affects residents	Use human review, monitoring, appeal routes and risk-tier-specific testing.
Technology is deployed because it is novel	Require a resident outcome, baseline and comparison to non-AI alternatives.

Review and renewal

AI and robotics rules will be reviewed at least annually and after material incidents, legal changes or evidence that a system creates unacceptable risk or fails to deliver public value.

Accountability and Public Reporting

Performance framework

AI and robotics reporting will track registered systems, risk reviews, incidents, resident complaints/appeals, pilot performance, workforce effects, compute use and any deployments stopped or redesigned.

Metric	Baseline	Target / direction	Horizon	Responsible body	Reporting
AI/automated systems in public registry	Establish initial inventory	100% of designated systems	First term	City Manager / technology / divisions	Quarterly
Consequential systems with impact assessment	Establish baseline	100% of designated systems	Ongoing	Program owner / privacy / legal	Per deployment
Resident cases with	Establish baseline	100% for covered services	First term	Service owners	Quarterly

Metric	Baseline	Target / direction	Horizon	Responsible body	Reporting
human escalation route					
Pilot outcome and stop-rule compliance	Pilot baseline	Publish every designated pilot result	Ongoing	Pilot owners	Per pilot
Workers receiving transition/training support before scale	Establish baseline	100% of affected designated pilots	Ongoing	HR / labour / program owner	Per pilot

Sources and Notes

Primary public sources are listed by organization and title. Campaign commitments, targets, scenarios and long-range aspirations are distinguished from adopted government policy throughout the report.

- [City of Toronto - Responsible Artificial Intelligence](#)
- [Office of the Privacy Commissioner of Canada](#)