

Specimen — illustrative example

Town of Quarry Bend is a fictional organisation. This document comes from the scenario described on riskwoven.com/examples, and is not an assessment of any real entity. It is advisory output to support your own decisions: not legal advice, not a certification, and not an audit.

Written and reviewed by RiskWoven.com. Written by RiskWoven.com's AI drafting agent and checked by RiskWoven.com's independent review stage before release. The review is automated. It is not a human, legal or audit review.

AI Governance and Accountability Policy

Town of Quarry Bend

Document type	policy	Version	1.0 Draft
Owner	Director of Corporate Services	Customer decision role	Chief Administrative Officer
Effective date	Upon customer release	Review cycle	Annual
Framework basis	RiskWoven AI practical best practices		
Jurisdiction	Ontario, Canada		
Prepared for	Town of Quarry Bend	Industry	Municipal government
Organization size	101-500	Locations	Not specified

Advisory draft for accountable-owner and professional review. Validate legal, regulatory, contractual, technical and operational requirements before customer release. This is not legal advice, a certification, or a compliance determination.

1. Purpose

This policy establishes the Town of Quarry Bend's mandate for governing artificial intelligence systems throughout their lifecycle. It ensures that AI systems serving residents and supporting

municipal operations are inventoried, assessed for risk, subject to appropriate human oversight, monitored for performance and incidents, and retired securely when no longer required.

2. Scope

This policy applies to all AI systems acquired, developed, operated or decommissioned by the Town of Quarry Bend, whether hosted on municipal infrastructure or provided by third parties. It covers systems used by all departments, contractors and service providers acting on behalf of the Town.

AI systems currently in scope include the automated triage system for 311 service requests and the first-pass screening tool for community grant applications. Any future AI system falls within scope from the point a department proposes its acquisition or development.

3. Definitions

AI system: Any software that uses machine learning, statistical modelling or rule-based logic to generate predictions, recommendations, classifications or decisions that influence municipal services or operations.

System owner: The director or manager accountable for an AI system's purpose, performance, compliance and lifecycle decisions. The system owner is recorded in the AI inventory.

Risk tier: A classification assigned to each AI system based on its potential impact on residents, equity of access, privacy, safety and municipal operations. The Town uses three tiers: High, Moderate and Low.

Impact assessment: A documented evaluation of an AI system's potential effects on residents, equity, privacy, accuracy and operational continuity, completed before deployment or material change.

Material change: Any alteration to an AI system's purpose, data inputs, decision logic, output use, supplier, hosting location or integration that could affect its risk profile or the rights of residents.

4. Policy Statements

4.1 The Town of Quarry Bend will maintain a current inventory of all AI systems, recording for each system its name, purpose, system owner, data processed, risk tier, deployment date, supplier (if any) and review status.

4.2 No AI system may be deployed or materially changed without a proportionate impact assessment reviewed by the appropriate authority for its risk tier.

4.3 Decision rights, human oversight requirements and escalation paths will be assigned to each AI system based on its risk tier and potential impact on residents.

4.4 AI systems will be monitored for performance degradation, incidents, complaints, model drift and supplier changes throughout their operational life.

4.5 AI systems will be retired securely when no longer required, with evidence preserved in accordance with the Records Retention and Disposal Policy.

4.6 AI systems may not be used for enforcement, licensing or any decision that directly determines a person's legal entitlement without explicit Council direction and additional safeguards.

5. AI Inventory

The Director of Corporate Services is accountable for maintaining the AI inventory. Each entry must include:

- System name and description
- Business purpose and intended outcomes
- System owner (by role)
- Data categories processed, referencing the Data Classification and Handling Standard
- Risk tier and date of last risk assessment
- Supplier name and contract reference (for third-party systems)
- Deployment date and current status
- Human oversight arrangements
- Date of next scheduled review

Departments must notify the Director of Corporate Services before acquiring, developing or proposing any AI system so that it can be added to the inventory. A system not recorded in the inventory may not be deployed.

6. Risk Classification

Each AI system is assigned a risk tier based on its potential impact. The assessment considers:

- Impact on residents' access to services or entitlements
- Equity implications, including differential effects on vulnerable groups
- Volume and sensitivity of personal information processed
- Consequences of erroneous outputs
- Reversibility of decisions influenced by the system
- Operational dependency and continuity risk

Risk tiers are defined as follows:

High: Systems whose outputs could materially affect a resident's access to services, entitlements or opportunities, or that process sensitive personal information at scale. The grant application screening tool is provisionally classified High because its outputs influence which applications proceed to committee review; the Town must confirm this assignment against the criteria in the AI Risk Management Standard.

Moderate: Systems that support operational decisions but whose outputs are routinely verified before action, and where errors are readily corrected. The 311 triage system is provisionally classified Moderate because misroutes are identified and corrected by receiving departments; the Town must confirm this assignment against the criteria in the AI Risk Management Standard.

Low: Systems with limited operational impact, no direct effect on residents' rights and minimal personal information processing.

Risk tier assignments are recorded in the AI inventory and reviewed at least annually or upon material change. The AI Risk Management Standard specifies the assessment method and criteria.

7. Impact Assessment and Review Before Deployment

Before any AI system is deployed or materially changed, the system owner must complete an impact assessment proportionate to the system's risk tier.

The assessment must address:

- The problem the system is intended to solve and the alternatives considered
- Data inputs, sources and any personal information involved
- How outputs will be used and by whom
- Potential for bias, error or inequitable outcomes
- Human oversight and intervention points
- Security, privacy and data residency requirements
- Supplier obligations, including those in the Third-Party Risk Management Policy
- Contingency if the system fails or is withdrawn

Review authority is assigned by risk tier:

- High-risk systems: Chief Administrative Officer, with a summary reported to Council
- Moderate-risk systems: Director of Corporate Services
- Low-risk systems: System owner, with notification to the Director of Corporate Services

No deployment may proceed until the review authority has recorded their review decision. The review record, including any conditions, is retained with the impact assessment.

8. Human Oversight and Decision Rights

Human oversight requirements are assigned based on risk tier and the nature of the AI system's outputs.

For High-risk systems:

- A designated staff member must review every output before it influences a decision affecting a resident.
- Any output that would exclude, deprioritise or disadvantage a resident must be reviewed by a person before action is taken.
- The system owner must designate a named role responsible for ongoing oversight.

For the grant application screening tool, a staff member confirms every shortlist before it proceeds to committee, and any application the model ranks low is read by a person before it is set aside.

For Moderate-risk systems:

- Outputs may be acted upon without individual review, provided a process exists to detect and correct errors promptly.
- The system owner must monitor error rates and escalate patterns to the Director of Corporate Services.

For the 311 triage system, the receiving department corrects misroutes and reports recurring patterns to the system owner.

For Low-risk systems:

- Standard operational oversight applies. The system owner monitors for anomalies.

Escalation: Any staff member who identifies an AI output that appears biased, erroneous or harmful must report it to the system owner. The system owner escalates to the Director of Corporate Services. Issues affecting residents' rights or safety are escalated to the Chief Administrative Officer.

9. Prohibited Uses

AI systems operated by or on behalf of the Town of Quarry Bend may not be used for:

- Enforcement actions, including by-law enforcement or penalty assessment
- Licensing decisions
- Any decision that directly determines a person's legal entitlement to a municipal service or benefit, unless Council has explicitly authorised the use and additional safeguards are in place
- Profiling residents for purposes unrelated to the service they have requested
- Any purpose that contravenes applicable Ontario privacy law or the Privacy and Personal Information Policy

Vendor models may not be trained on municipal data, and no resident information may leave Canada. These requirements are stated in the AI Acceptable Use Policy and enforced through contract terms specified in the Third-Party Risk Management Policy.

10. Monitoring and Ongoing Assurance

System owners are accountable for monitoring their AI systems throughout the operational lifecycle. Monitoring must address:

- Performance against intended outcomes
- Accuracy, error rates and false positive or negative patterns
- Complaints or concerns raised by residents or staff
- Evidence of model drift or degradation
- Supplier notifications of changes to the system, subprocessors or hosting
- Security incidents, in coordination with the AI Incident Response Procedure and Security Incident Response Procedure

High-risk systems require documented performance reviews on a cycle to be set and recorded by the Town in the AI Risk Management Standard. Moderate-risk systems require reviews at least annually. Low-risk systems are reviewed as part of the annual policy review or upon reported issues.

The Director of Corporate Services consolidates monitoring findings and reports to the Chief Administrative Officer on the metrics specified in Section 14.

11. Incident and Complaint Handling

AI-related incidents—including outputs that cause harm, privacy breaches, security events and significant errors—are managed under the AI Incident Response Procedure. Privacy breaches follow the Privacy Breach Response Procedure. Security incidents follow the Security Incident Response Procedure.

Complaints from residents or staff about AI system behaviour are recorded by the system owner, investigated and reported to the Director of Corporate Services. Patterns of complaints trigger a review of the system's risk tier and oversight arrangements.

The Chief Administrative Officer is notified of any incident involving a High-risk system or any incident that affects residents' rights, safety or trust in municipal services.

12. Supplier and Third-Party Obligations

AI systems provided by third parties are subject to the Third-Party Risk Management Policy and the Third-Party AI Review Procedure.

Contracts for AI systems must include the Town of Quarry Bend's required terms, as established by the Town and reflected in the Third-Party Risk Management Policy:

- Canadian data residency, with no processing of resident information outside Canada
- Prohibition on using municipal data to train vendor models
- Breach notification within the period the Town has set as its contractual standard (the Town has indicated 24 hours as its chosen requirement; this should be confirmed and recorded in the Third-Party Risk Management Policy and contract templates)
- Disclosure of subprocessors and advance notice of changes
- Cooperation with the Town's records retention obligations
- Return or secure destruction of municipal data at contract end

Supplier changes—including changes to the model, hosting, subprocessors or terms—must be reported to the system owner, who assesses whether the change is material and triggers a new impact assessment.

13. Retirement and Evidence Preservation

When an AI system is no longer required, the system owner initiates retirement in coordination with the Director of Corporate Services.

Retirement must include:

- Confirmation that the system is removed from operation and cannot process new data
- Secure return or destruction of data held by any third-party supplier, with written confirmation
- Preservation of records required under the Records Retention and Disposal Policy, including the impact assessment, review records, monitoring logs and incident reports
- Update of the AI inventory to reflect the system's retired status and retirement date

Retention periods for AI governance records are set by the Records Retention and Disposal Policy. Where that policy does not specify a period for a particular record type, the Director of Corporate Services determines the period based on legal, operational and accountability requirements.

14. Metrics and Reporting

The Director of Corporate Services reports to the Chief Administrative Officer on the following metrics at least annually:

- Number of AI systems in the inventory with a current risk assessment (target: 100%)
- High-impact AI uses with a designated human oversight role (target: 100%)
- Percentage of required reviews completed on time
- Number of approved exceptions past their expiry date (target: zero)
- Number of material control failures and overdue corrective actions

The Chief Administrative Officer reports a summary to Council annually, or more frequently if a significant incident or risk acceptance decision warrants it.

15. Exceptions

Exceptions to this policy may be granted only by the Chief Administrative Officer. Each exception must be documented, time-limited and include:

- The specific requirement being excepted
- The business justification
- Compensating controls in place
- Expiry date
- Conditions for renewal or closure

Exceptions are recorded in the AI inventory and tracked as part of the metrics in Section 14. An exception that remains open past its expiry date is reported to the Chief Administrative Officer for renewal or closure.

16. Roles and Responsibilities

Chief Administrative Officer: Accepts high residual risk. Reviews and decides on High-risk AI deployments. Grants exceptions. Reports to Council on AI governance.

Director of Corporate Services: Owns this policy. Maintains the AI inventory. Reviews Moderate-risk deployments. Consolidates monitoring and metrics. Coordinates retirement and evidence preservation.

System owners (directors and managers): Accountable for their AI systems' compliance with this policy. Complete impact assessments. Assign and maintain human oversight. Monitor performance and incidents. Initiate retirement.

All staff: Comply with the AI Acceptable Use Policy. Report concerns about AI outputs to the system owner.

17. Related Documents

- AI Acceptable Use Policy
- AI Risk Management Standard
- Third-Party Risk Management Policy
- Third-Party AI Review Procedure
- AI Incident Response Procedure
- Privacy and Personal Information Policy

- Privacy Breach Response Procedure
- Security Incident Response Procedure
- Data Classification and Handling Standard
- Records Retention and Disposal Policy
- Information Security Policy
- Access Control Policy
- Business Continuity Policy

18. Review and Maintenance

This policy is reviewed annually by the Director of Corporate Services, or sooner if triggered by:

- A significant AI-related incident
- Material changes to applicable Ontario privacy law or municipal obligations
- Introduction of a new High-risk AI system
- Direction from Council or the Chief Administrative Officer

Amendments are reviewed by the Chief Administrative Officer before taking effect.

19. Effective Date

This policy takes effect upon review by the Chief Administrative Officer and remains in force until superseded or withdrawn.

Owner review items

The drafting process identified the following in the supplied answers. Resolve each before customer release.

Type	Detail	Where
gap	The information supplied does not include a verified framework with specific control identifiers. This policy is drafted against RiskWoven AI practical best practices as described in the specification, but no framework clause numbers or citations are available to reference.	Entire document
gap	The information supplied does not specify the Town's existing risk-tier definitions or thresholds in detail. The three-tier model (High, Moderate, Low) and the provisional classification of the two current AI systems are inferred from the customer's answers. The AI Risk Management Standard should contain the detailed criteria; this policy defers to it. The Town must confirm both tier assignments.	Section 6 (Risk Classification)
gap	Retention periods for AI governance records are not specified in the information supplied. The policy defers to the Records Retention and Disposal Policy. The Town must confirm that policy covers impact assessments, review records, monitoring logs and incident reports, or must update it accordingly.	Section 13 (Retirement and Evidence Preservation)
vague	The customer's answers indicate 24-hour breach notification as the Town's chosen contractual requirement. This is recorded in Section 12 as the Town's own parameter, not as a framework or statutory mandate, and the Town is directed to confirm and record it in the Third-Party Risk Management Policy and contract templates.	Section 12 (Supplier and Third-Party Obligations)
gap	The information supplied references applicable Ontario privacy law but does not name the specific statutes. The policy refers to 'applicable Ontario privacy law' without enumerating instruments. Legal review should confirm which of MFIPPA, PHIPA or other provincial legislation applies to the Town's AI processing activities.	Section 9 (Prohibited Uses)
gap	The review frequency for High-risk systems is stated as a cycle to be set and recorded by the Town in the AI Risk Management Standard, because the information supplied does not mandate a specific interval. Quarterly is a reasonable starting point for the Town to consider. The Town must set this frequency and record it in the AI Risk Management Standard.	Section 10 (Monitoring and Ongoing Assurance)

Information the organization must supply

These values were not available from the intake and are required before this document can operate:

- Set and record the mandatory review frequency for High-risk AI systems in the AI Risk Management Standard (quarterly is a reasonable starting point for the Town to consider).
- Confirm that the Records Retention and Disposal Policy specifies retention periods for AI impact assessments, review records, monitoring logs and incident reports, or update that policy to include them.
- Confirm which Ontario privacy statutes apply to the Town's AI processing activities (e.g., MFIPPA, PHIPA or sector-specific legislation) and ensure the Privacy and Personal Information Policy reflects them.
- Confirm that 24-hour breach notification is the Town's chosen contractual standard and ensure it is recorded in the Third-Party Risk Management Policy and contract templates.
- Review the provisional risk-tier assignments for the 311 triage system (Moderate) and grant screening tool (High) and confirm they align with the criteria in the AI Risk Management Standard.
- Ensure the AI Risk Management Standard contains the detailed assessment criteria and scoring method referenced by this policy.

Written and reviewed by RiskWoven.com. Written by RiskWoven.com's AI drafting agent and checked by RiskWoven.com's independent review stage before release. The review is automated. It is not a human, legal or audit review.