CANADIAN SMALL BUSINESS AI ASSOCIATION

WHITE PAPER

GOVERNANCE & COMPLIANCE

# Responsible AI Use in Regulated Canadian Industries: A Practitioner Framework

A practical governance framework for small and mid-sized organizations deploying generative AI under Canadian privacy, safety, and professional conduct obligations — covering prompt design standards, output review requirements, staff accountability structures, and incident reporting practices.

ABOUT CSBAA

The Canadian Small Business AI Association (CSBAA) is a non-profit practitioner network supporting responsible AI adoption among Canadian organizations with 10–500 employees. This paper reflects consultation with member organizations in healthcare, financial services, legal services, and transportation between September 2025 and March 2026.

SUGGESTED CITATION

CSBAA. *Responsible AI Use in Regulated Canadian Industries: A Practitioner Framework.* White Paper CSBAA-WP-2026-03. Toronto: Canadian Small Business AI Association.

DISCLAIMER

This document is provided for general information and does not constitute legal advice. Organizations should obtain advice from qualified counsel on the application of federal, provincial, and sector-specific obligations to their circumstances.

## Executive summary

Generative AI has moved from experimentation to routine use in Canadian small and mid-sized organizations faster than governance has followed. In the CSBAA Practitioner Survey (March 2026, n = 412 organizations with 10–500 employees), 68% of respondents reported that staff use generative AI tools at least weekly — yet only 31% had a written AI use policy, and only 14% could name an individual accountable for AI outputs. In regulated settings, this gap is not merely an operational risk: it is a compliance exposure under PIPEDA, PHIPA, Quebec's Law 25, sector supervisory expectations, and professional conduct rules.

This white paper presents the CSBAA Responsible AI Use Framework: five principles translated into operating standards that a small compliance function can actually run. It specifies who is accountable for AI outputs (Section 3), what may and may not enter a prompt (Section 4), how much human review an output needs before it is relied upon (Section 5), and how AI-related incidents are recorded and escalated (Section 6). Section 7 provides a 90-day implementation roadmap and a four-level maturity model.

**68%**
of surveyed SMBs report weekly staff use of generative AI

**31%**
have a written AI use policy in force

**22%**
experienced at least one AI-related data-handling near-miss in the past year

**9%**
have a defined AI incident reporting procedure

Source: CSBAA Practitioner Survey, March 2026 (n = 412). See Reference [10].

# Contents



## 1 The regulatory landscape

No Canadian statute currently regulates generative AI as such. The proposed Artificial Intelligence and Data Act (AIDA), introduced as part of Bill C-27, died on the Order Paper when Parliament was prorogued in January 2025, and no successor bill had been tabled at the time of writing. In practice, this means AI use by Canadian SMBs is governed by the *existing* body of privacy, safety, and professional conduct law — applied to a new class of tool. Three layers matter.

### 1.1 Federal and provincial privacy law

PIPEDA applies to personal information handled in the course of commercial activity in most provinces. Entering a customer's personal information into a third-party AI tool is a *disclosure* to that provider: it engages PIPEDA's consent, limiting-use, and safeguards principles, and the accountability principle makes the organization — not the vendor — answerable for what happens next. The Office of the Privacy Commissioner's 2023 principles for generative AI make clear that the regulator expects organizations to apply established privacy principles to these tools now, without waiting for AI-specific legislation.

Provincial regimes add sharper edges. Ontario's PHIPA governs personal health information held by health information custodians; disclosing PHI to an AI vendor that is not an agent or contracted service provider under an appropriate agreement can constitute an unauthorized disclosure, and PHIPA's mandatory breach notification applies. Quebec's Law 25 goes further than PIPEDA: it requires privacy impact assessments before communicating personal information outside Quebec, and grants individuals the right to be informed when a decision about them is based exclusively on automated processing.

### 1.2 Sector supervisory expectations

Regulated sectors carry obligations that reach AI use directly, even where the word "AI" never appears in the statute. The table below summarizes the instruments CSBAA members most often need to map their AI use against.

| SECTOR | KEY INSTRUMENTS | WHAT THEY MEAN FOR AI USE |
|---|---|---|
| **Healthcare** | PHIPA (ON) and provincial equivalents; college standards of practice (e.g., CPSO) | PHI may not enter tools outside the custodian's agreements; clinicians remain professionally responsible for AI-assisted documentation and cannot delegate clinical judgment. |
| **Financial services** | OSFI Guideline E-23 (model risk, effective May 2027); FSRA / provincial licensing rules; PIPEDA | AI models used in decisions require inventory, validation, and lifecycle governance; suitability and fair-treatment obligations apply to AI-assisted advice. |
| **Legal services** | Law society rules of professional conduct; LSO generative AI guidance (2024); court practice directions | Competence and confidentiality duties cover AI use; several courts require disclosure or certification of AI-assisted filings; hallucinated citations have led to costs awards. |
| **Transportation** | Transport Canada safety management system (SMS) requirements; provincial carrier safety regimes | AI-assisted scheduling, maintenance, or safety documentation falls inside the SMS; safety-critical decisions require documented human sign-off. |

Table 1 — Sector instruments most relevant to generative AI use by Canadian SMBs.

### 1.3 Voluntary standards

ISO/IEC 42001:2023 (AI management systems) and the NIST AI Risk Management Framework are increasingly cited in procurement questionnaires and insurance underwriting. Neither is mandatory for Canadian SMBs, but the CSBAA Framework is deliberately mappable to both, so that an organization adopting it can answer such questionnaires without a parallel exercise.

# 2 The CSBAA Responsible AI Use Framework

The Framework rests on five principles. Each principle is expressed as a standard an organization can adopt verbatim in its AI use policy, and each is operationalized by one of the sections that follow. The design constraint throughout is that the framework must be runnable by an organization with no dedicated AI team — typically a compliance lead, an operations manager, and department supervisors.

**P1**
**Lawful & documented purpose**
Every approved AI use case is recorded with its purpose and legal basis.

**P2**
**Data minimisation by default**
Prompts contain the least identifying data that achieves the task. (§4)

**P3**
**Named human accountability**
A named person owns every AI output relied upon. (§3)

**P4**
**Proportionate output review**
Review effort scales with the consequence of the output. (§5)

**P5**
**Transparent incident handling**
AI-related incidents are logged, classified, and escalated. (§6)

Figure 1 — The five principles of the CSBAA Responsible AI Use Framework.

## 2.1 Principle statements

**P1 — Lawful and documented purpose.** The organization maintains a register of approved AI use cases. Each entry records the tool, the purpose, the categories of data involved, the legal basis for any personal information processing, and the responsible role. Staff use of AI outside a registered use case requires prior approval from the AI Use Lead.

**P2 — Data minimisation by default.** Prompts are drafted to exclude personal information, client identifiers, and confidential business information unless the registered use case expressly permits them and the tool is under an appropriate agreement. De-identification is the default, not the exception.

**P3 — Named human accountability.** Every AI output that leaves the organization or informs a decision has a named human owner who has reviewed it to the standard Section 5 requires. "The AI produced it" is never an account of responsibility.

**P4 — Proportionate output review.** Review requirements are set by the consequence of the output, not by the effort of producing it. A three-tier review scheme (Section 5) distinguishes internal drafts, client-facing material, and regulated determinations.

**P5 — Transparent incident handling.** AI-related incidents and near-misses are recorded in the same discipline as other operational incidents: a log, a severity classification, defined escalation, and a link into statutory breach notification where personal information is involved.

**CASE SCENARIO 1 — COMMUNITY HEALTH CLINIC, ONTARIO (14 STAFF)**

An intake coordinator uses a general-purpose chatbot to "tidy up" referral summaries, pasting in full referral letters including patient names and health card numbers. Under the Framework, this fails P1 (no registered use case), P2 (identifiable PHI in a prompt to an uncontracted tool), and constitutes a Severity 2 incident under Section 6 — a potential unauthorized disclosure under PHIPA requiring assessment for notification. The corrective path: register a documentation-support use case on a tool covered by an agreement with the custodian, adopt the de-identification pattern in Section 4.3, and train intake staff on the prompt standards.

## 3 Role-based accountability

Accountability fails when it is collective. The Framework defines four roles; in a small organization one person may hold two of them, but no role may be vacant, and the Accountable Executive and Reviewer roles may not be held by the same person for Tier 3 outputs (Section 5). Only 14% of surveyed organizations could name an accountable individual for AI outputs; this section is designed to make that assignment a one-page exercise.

| ROLE | TYPICALLY HELD BY | CORE RESPONSIBILITIES |
|---|---|---|
| **Accountable Executive** | Owner, managing partner, executive director, or CEO | Approves the AI use policy and use-case register; owns residual risk; receives Severity 1–2 incident escalations; reports to the board or partnership. |
| **AI Use Lead** | Compliance lead, privacy officer, or operations manager | Maintains the use-case register and tool list; approves new use cases; runs training; administers the incident log; conducts the annual review. |
| **Reviewer** | Supervisor or licensed professional in the relevant function | Performs Tier 2–3 output review (Section 5); signs off before reliance; is the professional of record where a licence attaches to the work. |
| **Operator** | Any staff member using an AI tool | Uses only registered tools and use cases; follows prompt standards (Section 4); performs Tier 1 self-review; reports incidents and near-misses without penalty. |

Table 2 — The four Framework roles.

### 3.1 The professional-of-record rule

Where the output is work that a licence, designation, or statutory duty attaches to — a clinical note, a legal filing, a suitability assessment, a safety inspection record — the Reviewer must be the professional of record, and the review obligation cannot be discharged by anyone else, however senior. Professional regulators have been consistent on this point: AI assistance does not dilute the licensee's personal responsibility for the work product.

The corollary is a no-blame reporting rule for Operators. An Operator who reports their own near-miss — a prompt that should not have been sent, an output relied on without review — is not disciplined for the report itself. Survey data suggest why this matters: of the 22% of organizations reporting a near-miss, fewer than half learned of it from the person involved.

## 4 Prompt design standards

Most AI compliance failures in SMBs happen at the prompt, not the model: information enters a tool it should never have reached. The Framework treats a prompt to an external AI tool as a disclosure to the tool's provider and classifies prompt content in three bands.

| BAND | CONTENT | RULE |
|---|---|---|
| Green | Public information, generic drafting, de-identified patterns, published law and guidance, the organization's own published material | Permitted in any registered tool. |
| Amber | Confidential business information; personal information that has been de-identified but could be re-identified in context | Only in tools with a contract covering confidentiality and no training on inputs; only within a registered use case. |
| Red | Identifiable personal information, PHI, financial account data, privileged material, credentials, safety-critical operational data | Never in a general-purpose tool. Only where a specific registered use case, a data-processing agreement, and the AI Use Lead's written approval all exist. |

Table 3 — Prompt content bands.

### 4.1 What should never enter a prompt

- Names, addresses, health card or SIN numbers, dates of birth, or any combination that identifies a person, outside a Red-band approved use case.
- Material subject to solicitor-client privilege or litigation privilege.
- Authentication credentials, API keys, or system configuration details.
- Third-party confidential information received under NDA.
- Content the organization is contractually or statutorily required to keep within specified systems or jurisdictions.

### 4.2 What a good prompt contains instead

The task, the audience, the format, the constraints — and placeholders for the facts. A prompt that reads "Draft a letter declining a mortgage application because the debt-service ratio exceeds our threshold; neutral professional tone; applicant is [NAME]" produces output of the same quality as one containing the real file, and produces no disclosure. Staff should be trained to write the placeholder pattern first and treat any urge to paste source documents as the signal to check the band.

### 4.3 The de-identification pattern

Where source material is genuinely needed, the Framework's default pattern is: copy the source into a local working document; replace direct identifiers with bracketed tokens ([CLIENT-A], [DOB], [ACCOUNT]); review for indirect identifiers (rare conditions, unique job titles, small-town locations); submit the tokenized text; re-insert identifiers only in the organization's own systems. The Operator, not the tool, performs the substitution — and for Amber-band material the tokenized text still travels only to contracted tools.

**CASE SCENARIO 2 — MORTGAGE BROKERAGE, BRITISH COLUMBIA (23 STAFF)**

A broker drafts client-suitability narratives with an AI assistant, pasting bank statements into the prompt. Financial account data is Red band: the brokerage registers a document-drafting use case on a contracted tool, adopts the placeholder pattern for narratives, and prohibits statement uploads. Reviewer sign-off (the supervising broker) is required before any narrative enters a client file — aligning the workflow with both the prompt standards and provincial suitability obligations.

## 5 Output review requirements

Generative models produce fluent, confident text whether or not it is correct. The Framework therefore ties review effort to the consequence of relying on the output, in three tiers. The tier is set by where the output is going, not by how it was made: a paragraph destined for a regulated determination is Tier 3 even if the model only "polished the wording."

| TIER | OUTPUT DESTINATION | REQUIRED REVIEW | RECORD KEPT |
|---|---|---|---|
| **Tier 1** | Internal drafts, brainstorming, research starting points | Operator self-review; facts verified before the draft is reused | None |
| **Tier 2** | Client-facing communications, published content, internal records of substance | Reviewer verifies facts, figures, names, and every citation against source | Reviewer initials and date on the record |
| **Tier 3** | Regulated determinations: clinical documentation, filings, suitability assessments, safety records, decisions about individuals | Professional of record performs full substantive review and adopts the content as their own work | Signed attestation; AI assistance noted where a court or regulator requires disclosure |

Table 4 — The three review tiers.

### 5.1 The citation rule

Every authority, statistic, quotation, or reference in an AI-assisted output is verified against its original source before Tier 2 or Tier 3 reliance — no exceptions, including references that "look right." Canadian courts have now sanctioned lawyers in multiple matters for filing AI-fabricated citations; the pattern generalizes to any regulated document. If a source cannot be located, the claim it supports is removed.

### 5.2 Automated decisions about individuals

The Framework prohibits fully automated decisions with legal or similarly significant effects on individuals — hiring, credit, coverage, care — in all member deployments. A human with authority to change the outcome must make the final determination, and where Quebec's Law 25 applies, individuals must be informed of automated processing and given a channel to have the decision reviewed by a person.

## 6 Incident reporting and escalation

An AI incident is any event in which AI use caused, or could plausibly have caused, harm to a client, a breach of law or professional duty, or material error in the organization's records. The Framework classifies incidents in three severities and routes each through a fixed escalation path. The log itself is short — one row per event — but its existence changes behaviour: it converts "quiet fixes" into learning.

| SEVERITY | DEFINITION AND EXAMPLES | RESPONSE |
|---|---|---|
| **Severity 1** | Confirmed harm or confirmed breach: personal information disclosed to an uncontracted tool with real risk of significant harm; an erroneous AI-assisted determination acted upon; a fabricated authority filed | Escalate to Accountable Executive within 24 hours; assess statutory notification (PIPEDA s. 10.1, PHIPA, Law 25); notify affected parties and regulators as required; suspend the use case pending review. |
| **Severity 2** | Potential breach or contained error: Red-band content sent to a contracted tool outside its use case; a Tier 2 error caught after sending but corrected without harm | AI Use Lead investigates within 5 business days; documents cause and correction; reports pattern-level findings to the Accountable Executive quarterly. |
| **Severity 3** | Near-miss: a non-compliant prompt drafted but not sent; an error caught in review; use of an unregistered tool with Green-band content only | Logged, no investigation required; reviewed in aggregate at the annual policy review to identify training needs. |

Table 5 — Incident severity classification.

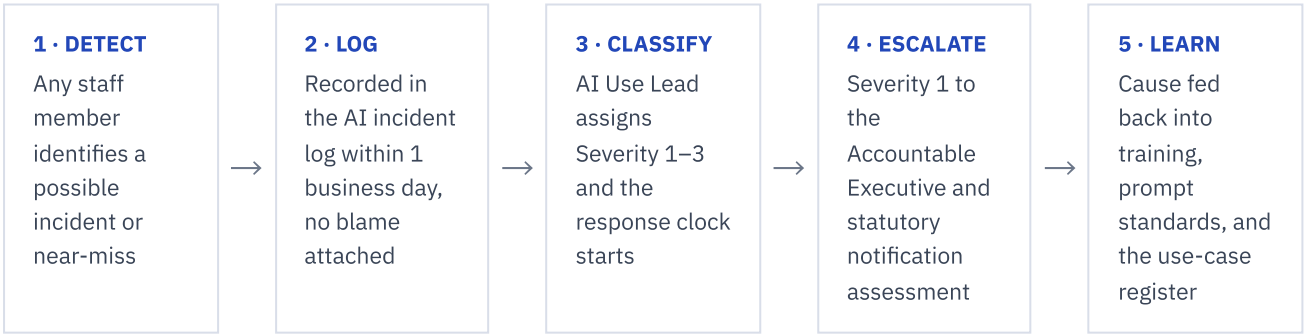


Figure 2 — The five-step incident path.

**CASE SCENARIO 3 — LITIGATION BOUTIQUE, ALBERTA (9 LAWYERS)**

An associate's AI-assisted brief includes two authorities that do not exist; the error is caught by the reviewing partner before filing. Under the Framework this is a Severity 3 near-miss — logged, no discipline — but the aggregate review reveals three similar catches in one quarter. The firm's response: the citation rule (§5.1) becomes a checklist item on every filing, and legal-research use cases are restricted to a tool with retrieval against a verified case-law database.

# 7 Implementation: a 90-day roadmap

The Framework is designed to be stood up in one quarter without external consultants. The sequence matters: inventory before policy, policy before training, training before enforcement.

**DAYS 1–30 · SEE**

- Appoint the four roles
- Survey actual AI use, amnesty declared
- Build the tool list and use-case register
- Identify Red-band exposures for immediate stop

**DAYS 31–60 · SET**

- Adopt the AI use policy (Appendix A)
- Contract review of retained tools
- Define review tiers per department
- Open the incident log

**DAYS 61–90 · RUN**

- Train all Operators on bands and tiers
- Reviewer sign-off live for Tier 2–3
- First monthly incident-log review
- Schedule the annual policy review

Figure 3 — The 90-day implementation sequence.

## 7.1 The CSBAA maturity model

Organizations self-assess annually against four levels. Level 2 is the minimum CSBAA considers defensible for a regulated SMB; most surveyed organizations currently sit at Level 0 or 1.

| LEVEL | CHARACTERISTICS |
|---|---|
| **0 — Unmanaged** | AI use is happening but invisible to management; no policy, no inventory, no accountable person. |
| **1 — Aware** | A policy exists, typically prohibitive; actual practice diverges from it; no review or incident discipline. |
| **2 — Managed** | Roles assigned; use-case register and prompt bands in force; Tier 2–3 review operating; incident log open. |
| **3 — Improving** | Incident-log patterns drive changes to standards and training; annual review documented; framework mapped to ISO/IEC 42001 or NIST AI RMF for procurement and insurance. |

Table 6 — CSBAA AI governance maturity model.

Governance of AI in regulated Canadian SMBs does not await new legislation; the obligations already exist, and the tools are already in use. The gap is operational, and it is closable in a quarter. Organizations that adopt the Framework report a second-order benefit

beyond compliance: staff use AI more, not less, once the rules of the road are explicit — because the fear of doing it wrong is replaced by a clear account of how to do it right.

## A Appendix A — AI use policy adoption checklist

A completed checklist constitutes prima facie evidence of Level 2 maturity. Items are ordered as they should be completed.

- [ ] **Roles.** Accountable Executive, AI Use Lead, Reviewers, and Operator population named in writing; no vacancies.
- [ ] **Tool list.** Every AI tool in actual use identified; each marked contracted / uncontracted; uncontracted tools restricted to Green band.
- [ ] **Use-case register.** Each approved use case records tool, purpose, data categories, legal basis, and responsible role.
- [ ] **Prompt standards.** Green/Amber/Red bands adopted; never-list (§4.1) circulated; de-identification pattern documented.
- [ ] **Review tiers.** Tier definitions mapped to the organization's actual document types; professional-of-record rule confirmed for licensed work.
- [ ] **Incident log.** Log opened; severity definitions adopted; escalation contacts posted; no-blame reporting rule communicated.
- [ ] **Training.** All Operators trained on bands, tiers, and reporting; completion recorded; new-hire onboarding updated.
- [ ] **Statutory mapping.** Breach-notification triggers (PIPEDA s. 10.1, PHIPA, Law 25 as applicable) linked from the incident procedure; privacy impact assessment completed where Law 25 requires one.
- [ ] **Annual review.** Date set; scope includes incident-log patterns, tool list currency, and regulatory change since last review.

## B Glossary

**Generative AI.** Systems that produce text, images, or other content from prompts, typically built on large language models.

**Prompt.** The full content submitted to an AI tool, including pasted documents and uploaded files, not only the typed instruction.

**Contracted tool.** An AI tool covered by an agreement addressing confidentiality, data residency where required, and non-training on the organization's inputs.

**De-identification.** Removing or tokenizing direct and indirect identifiers so an individual cannot reasonably be identified from the material.

**Human-in-the-loop.** A process design in which a person with authority to change the outcome reviews AI output before it takes effect.

**Near-miss.** An event that could have caused harm or breach but was caught before any effect occurred.

**PHI / PII.** Personal health information (as defined in PHIPA and equivalents) / personally identifiable information generally.

**Hallucination.** Fluent AI output that is factually false or cites sources that do not exist.

**Professional of record.** The licensed individual to whom professional responsibility for a work product attaches.

**Use-case register.** The organization's record of approved AI uses: tool, purpose, data categories, legal basis, responsible role.

## C References


1. *Personal Information Protection and Electronic Documents Act* (PIPEDA), S.C. 2000, c. 5.
2. *Personal Health Information Protection Act* (PHIPA), S.O. 2004, c. 3, Sched. A.
3. Quebec, *Act respecting the protection of personal information in the private sector*, CQLR c. P-39.1, as amended by the *Act to modernize legislative provisions as regards the protection of personal information* (Law 25), S.Q. 2021, c. 25.
4. Bill C-27, *Digital Charter Implementation Act, 2022*, 1st Sess., 44th Parl. (including the proposed *Artificial Intelligence and Data Act*); died on the Order Paper, January 2025.
5. Office of the Superintendent of Financial Institutions, *Guideline E-23: Model Risk Management*, final version September 2024, effective May 1, 2027.
6. Office of the Privacy Commissioner of Canada, *Principles for Responsible, Trustworthy and Privacy-Protective Generative AI Technologies*, December 2023.
7. Law Society of Ontario, *Licensee Guidance on the Use of Generative Artificial Intelligence*, April 2024.
8. ISO/IEC 42001:2023, *Information technology — Artificial intelligence — Management system*. Geneva: International Organization for Standardization.
9. National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, January 2023.
10. Canadian Small Business AI Association, *Practitioner Survey on AI Use and Governance*, fieldwork March 2026, n = 412 Canadian organizations (10–500 employees). Unpublished member survey; methodology available on request.

CSBAA-WP-2026-03 · Licensed CC BY-NC 4.0