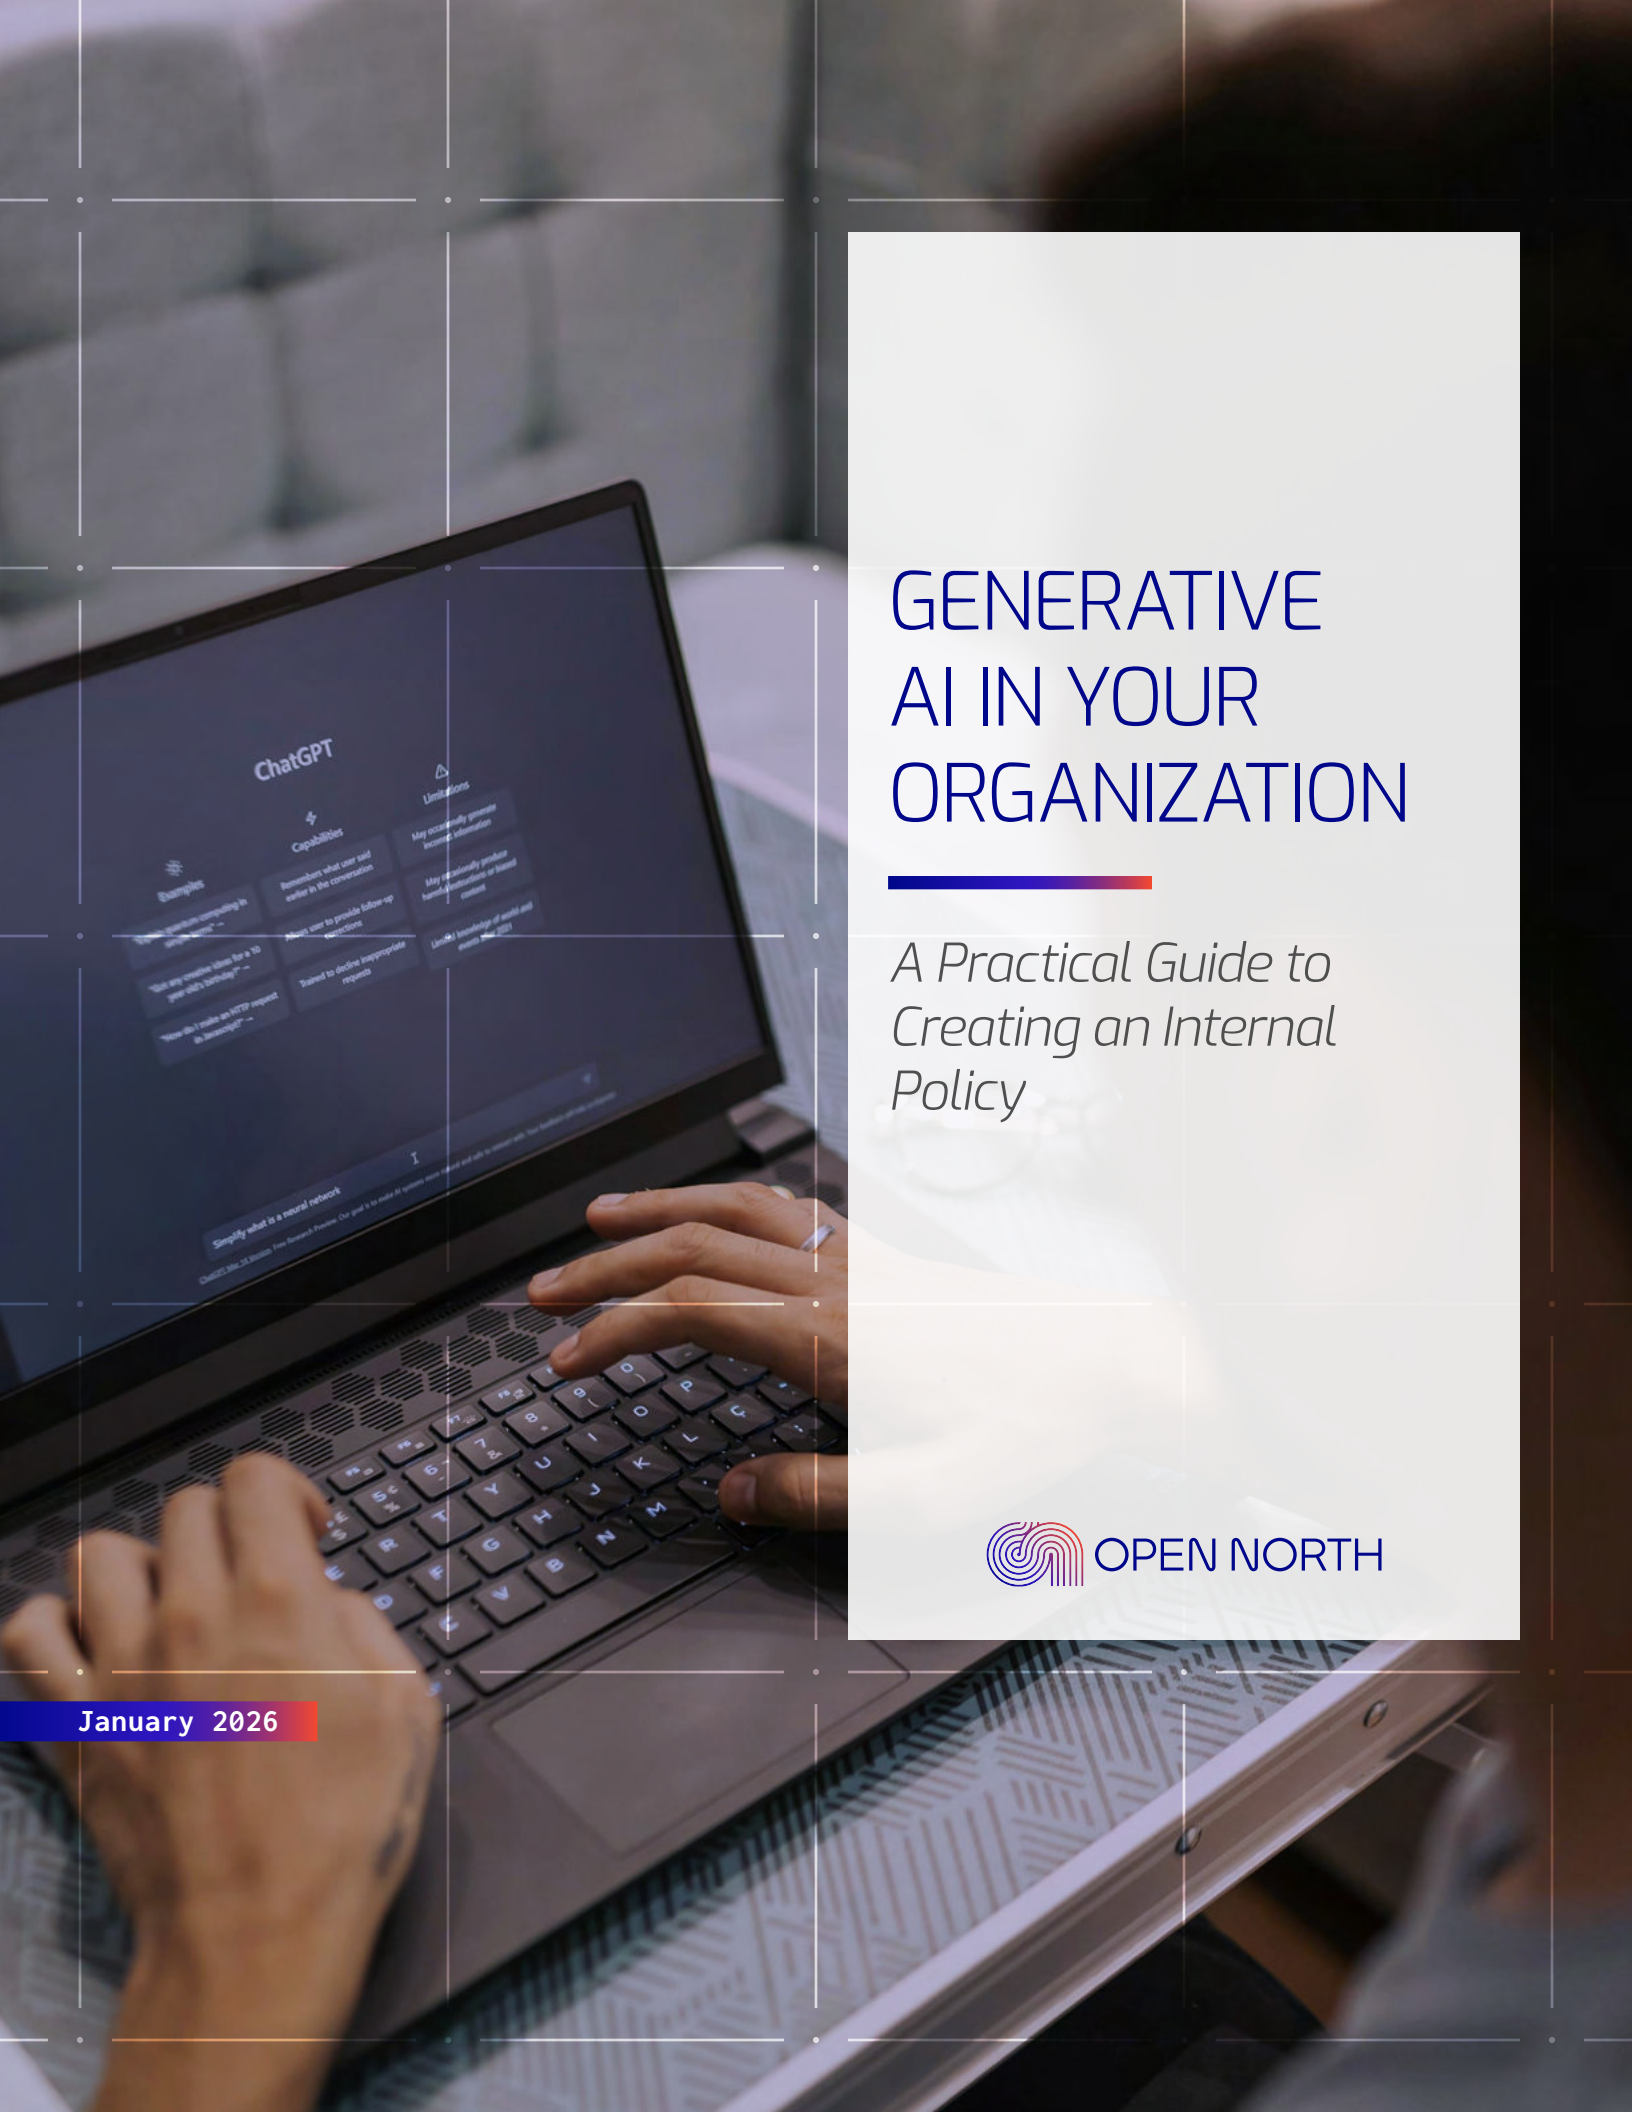

GENERATIVE AI IN YOUR ORGANIZATION

A Practical Guide to Creating an Internal Policy

OPEN NORTH

January 2026

ABOUT THIS GUIDE

This guide is designed for small and medium-sized organizations in Québec and across Canada. It provides both educational content to help you understand generative AI risks and opportunities, and practical, actionable rules you can adapt for your organization.

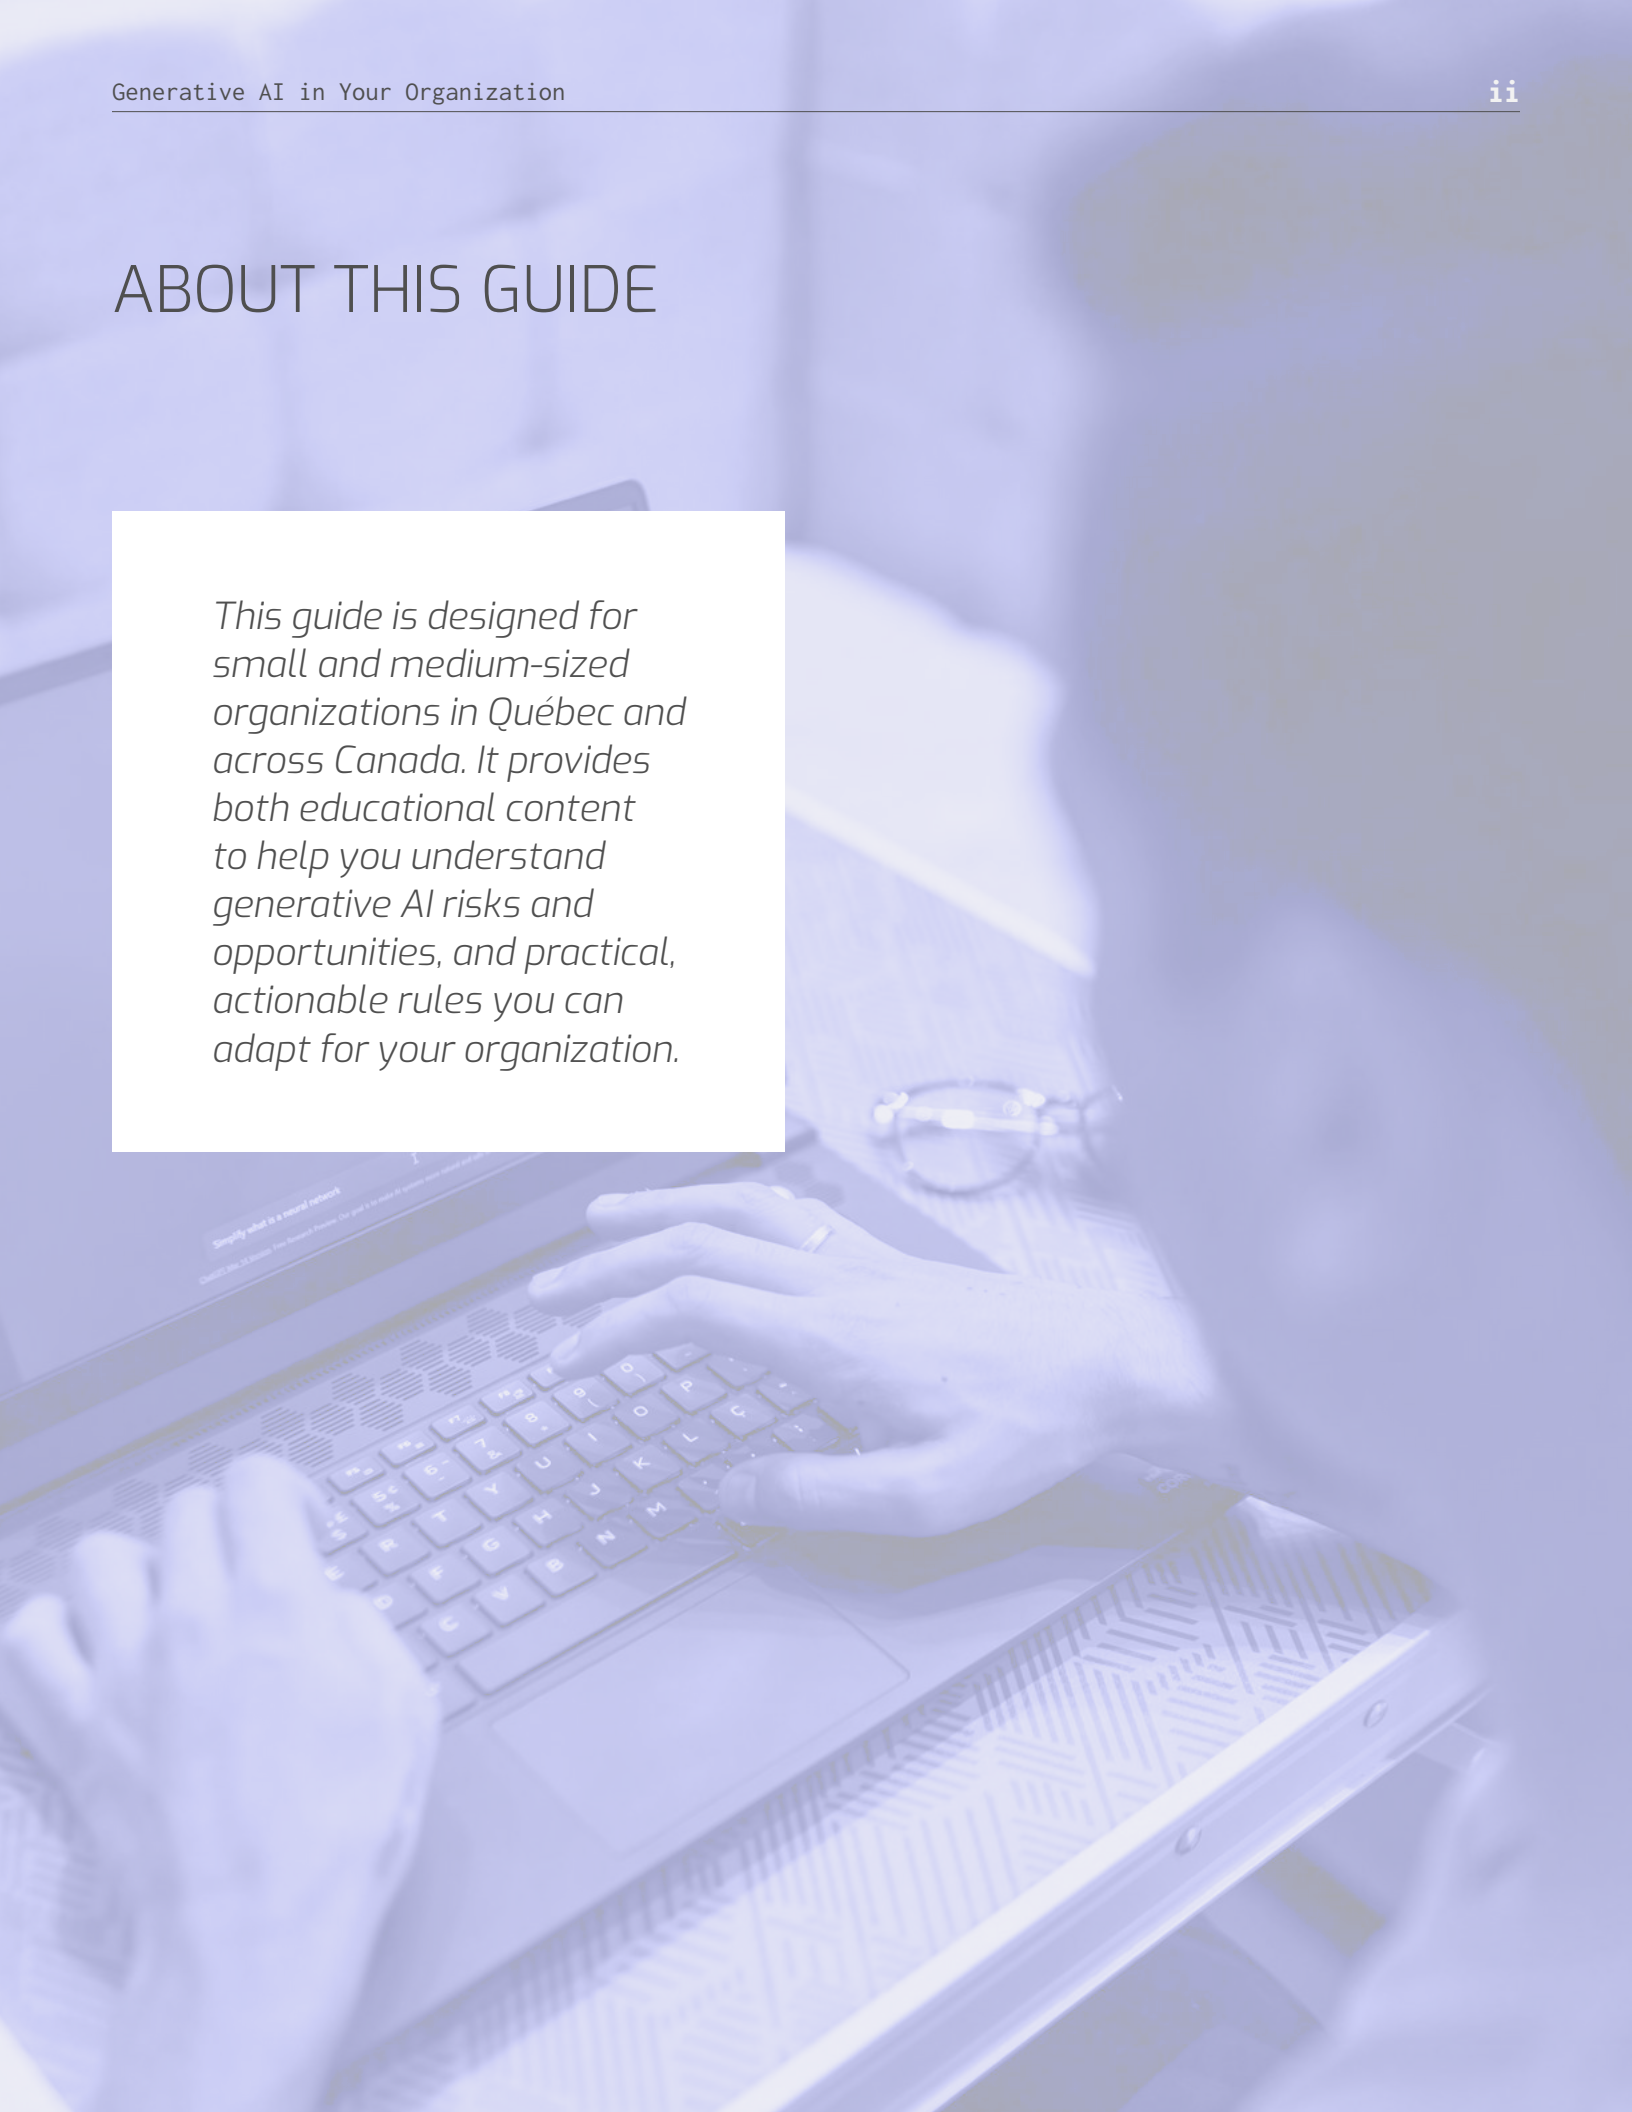

CONTENTS

WHY YOUR ORGANIZATION NEEDS A GENERATIVE AI POLICY

The Reality: AI Is Already in Your Workplace

Generative AI tools like ChatGPT, Microsoft Copilot, and Google Gemini have become part of everyday work. Some research sources indicate that approximately 75% of knowledge workers bring their own AI tools to work, and over 60% of those using AI are unaware of their organization's policies. This creates **"shadow AI"**: unmanaged use that exposes your organization to risks you cannot see or control.

Documented Risks: Real Cases, Real Consequences

- **Organizational Liability (Air Canada, 2024):** In Moffatt v. Air Canada, the BC Civil Resolution Tribunal held that Air Canada was liable for incorrect information provided by its AI chatbot. The tribunal rejected the argument that the chatbot was "a separate legal entity responsible for its own actions," stating it should be "obvious to Air Canada that it is responsible for all the information on its website." This case establishes that organizations cannot disclaim responsibility for AI outputs.
- **Data Exposure (Samsung, 2023):** Within 20 days of allowing ChatGPT access, Samsung experienced three separate data leaks—source code, program optimization code, and meeting transcripts. This data cannot be retrieved from the AI provider's servers once entered.
- **Legal Hallucinations:** Courts in Canada and globally have documented hundreds of cases where lawyers submitted AI-fabricated legal citations. In Zhang v. Chen (2024), a BC lawyer was ordered to pay costs for submitting fake AI-generated citations. In Reddy v. Saroya (2025), Alberta's Court of Appeal held lawyers bear "ultimate responsibility" for AI errors.
- **Bias and Discrimination:** In the iTutorGroup EEOC settlement (USA, 2023), an AI hiring tool automatically rejected female applicants 55+ and male applicants 60+, resulting in a $365,000 settlement.

> Key Takeaway
>
> Your organization is responsible for all AI outputs used on its behalf. "The AI made an error" is not a legal defence.

The Opportunity: Responsible AI Can Help

When used responsibly with proper safeguards, generative AI can genuinely help resource-constrained organizations: drafting routine communications, summarizing documents, supporting multilingual accessibility, exploring ideas, and assisting research. The goal of this guide is to help you capture these benefits while managing the real risks.

WHAT IS GENERATIVE AI? KEY CONCEPTS

Before you can govern something, you need to understand what it is. This section provides essential definitions focused on what staff need to know to use these tools safely.

Essential Definitions

TERM	DEFINITION (WHAT YOU NEED TO KNOW)
Generative AI	AI systems that create new content—text, images, code, audio—based on patterns learned from training data. Examples: ChatGPT, Copilot, Gemini, Claude.
Hallucination	When AI generates content that sounds plausible but is factually incorrect or fabricated. This is a fundamental characteristic of how these systems work—not a bug that will be fixed. Always verify AI outputs and look for sources.
Prompt	The input you give to an AI system. Never include personal information, confidential data, or trade secrets in prompts to public AI tools.
Personal Information	Any information about an identifiable individual: names, contact info, identifiers, health data, financial records, or any data that alone or combined could identify a person.
Shadow AI	Unauthorized use of AI tools by staff, often using personal accounts or unapproved tools. Creates compliance, security, and quality risks your organization cannot see or manage.

Types of AI Tools: Risk Profiles

TOOL TYPE	EXAMPLES	KEY RISKS
Public Cloud (Free)	ChatGPT Free, Claude Free, Gemini Free, Perplexity Free, Copilot Free	**HIGH RISK:** Data may be used for training; limited privacy controls; servers likely outside Canada. Never use for confidential or personal data
Public Cloud (Consumer Paid)	ChatGPT Plus or Pro, Claude Pro or Max, Gemini Pro or Ultra, Perplexity Pro or Max, Copilot Pro	**MEDIUM-HIGH RISK:** Opt-out from training usually available; some privacy controls; servers still likely outside Canada. Never use for confidential or personal data. Check their data training policies before use with internal/relevant data
Enterprise Cloud	ChatGPT Enterprise, Microsoft 365 Copilot, Google Workspace AI	**MEDIUM RISK:** Data typically not used for training; better privacy controls; requires PIA for personal data. Verify data residency and terms
Embedded AI	AI in Adobe, Canva, Zoom, browser extensions	**VARIABLE RISK:** Often enabled by default; review vendor terms before use. Browser AI extensions may read all content
Self-Hosted	Llama, Mistral, other open-source AI models on own infrastructure	**LOWER RISK:** Complete data control; no external sharing. Requires technical expertise

What Generative AI Cannot Do

- **Cannot guarantee accuracy:** AI predicts plausible text, not correct text. It will confidently state false information.
- **Cannot maintain confidentiality on public tools:** Data entered may be stored, reviewed, or used for training.
- **Cannot replace human judgment:** For decisions affecting individuals, human review is essential and often legally required.

LEGAL FRAMEWORK: QUÉBEC, CANADA, AND BEYOND

Legal Disclaimer

This section provides general information about applicable legal requirements. It is not legal advice and must be validated by your legal and compliance advisors. Laws and regulations may have changed since this guide was published.

Québec: Law 25 (Loi 25) — Now Fully in Effect

Law 25 became fully effective on September 22, 2024. It is currently the strictest privacy law in Canada. If your organization operates in Québec or handles data of Québec residents, you must comply. Key provisions affecting AI use:

- **Automated Decision-Making (Section 12.1):** When decisions are made exclusively by automated processing of personal information, you must inform the individual, explain the personal information used, the reasons and principal factors leading to the decision, and provide the right to submit observations to a person who can review the decision.
- **Privacy Impact Assessments (PIAs):** Required before acquiring or developing new technologies that process personal information, and before communicating personal information outside Québec. Most cloud AI tools constitute a transfer outside Québec.
- **Designated Privacy Officer:** By default, the highest-ranking person in the organization is responsible for privacy compliance. This can be delegated in writing.
- **Biometric Data:** Use of biometric identification systems requires notification to the Commission d'accès à l'information (CAI) at least 60 days before implementation.
- **Data Portability:** Individuals have the right to receive their personal information in a structured, commonly used format.

Penalties: Administrative fines up to $10 million CAD or 2% of worldwide turnover. Penal sanctions up to $25 million CAD or 4% of global revenue. Private right of action with minimum $1,000 CAD damages per individual.

Important CAI Guidance

The *Commission d'accès à l'information* has issued guidance that directly affects AI use:

- **Inferred Data as New Collection:** In its November 2022 decision (Val-des-Cerfs), the CAI determined that when AI generates predictions or inferences about individuals (e.g., predicting employee turnover risk or inferring customer income), this constitutes a new collection of personal information, triggering all applicable privacy requirements.
- **Meaningful Human Review:** For a decision not to be considered "exclusively automated," the human review must be substantive—not merely approving an AI suggestion without critical analysis.

Federal: Current Status

Canada does not currently have comprehensive federal AI-specific legislation. Bill C-27, which included the proposed Artificial Intelligence and Data Act (AIDA), died when Parliament was prorogued on January 6, 2025. However:

- **PIPEDA still applies:** The Personal Information Protection and Electronic Documents Act applies to private-sector commercial activities federally regulated or in provinces without substantially similar legislation.
- **Joint Privacy Commissioners' Principles:** In December 2023, all federal, provincial, and territorial privacy commissioners issued joint principles for generative AI. While not technically binding, these represent authoritative guidance on compliance expectations.
- **Voluntary Code of Conduct:** The federal government maintains a voluntary code of conduct for advanced generative AI systems (September 2023).

Liability Even Without AI-Specific Laws

Even in the absence of specific AI legislation, liability for AI-related harms can arise through:

- **Civil liability / tort law:** Negligence, negligent misrepresentation (as in the Air Canada case)
- **Contractual obligations:** Breach of confidentiality, failure to meet service standards
- **Privacy law:** Violations of Law 25, PIPEDA, or provincial privacy statutes
- **Human rights law:** Discrimination resulting from biased AI outputs
- **Professional duties:** For regulated professionals (lawyers, accountants, healthcare providers)

International Frameworks (Reference)

While not directly binding in Canada, these frameworks inform best practices and may apply to organizations with international operations:

- **EU AI Act:** In force August 2024, full application by August 2026. Has extraterritorial reach affecting Canadian organizations serving EU customers.
- **NIST AI Risk Management Framework:** Comprehensive voluntary framework from the US; widely adopted internationally.
- **ISO/IEC 42001:2023:** First certifiable AI management system standard.

RISK CLASSIFICATION AND DATA PROTECTION

Data Classification for AI Inputs

Before using any AI tool, classify the data you intend to input. This determines what tools you can use and what approvals are required.

LEVEL	EXAMPLES	AI TOOL PERMITTED
Restricted	SINs, health records, credit cards, biometric data, children's data, passwords	**NEVER share with any AI tool**
Confidential	Trade secrets, financial records, legal matters, HR files, client case details	Enterprise tools with PIA completed; never public tools
Internal	Draft internal documents, meeting notes (anonymized), general procedures	Approved enterprise tools; public tools only if data is anonymized
Public	Published materials, public info, general research questions	Any approved tool

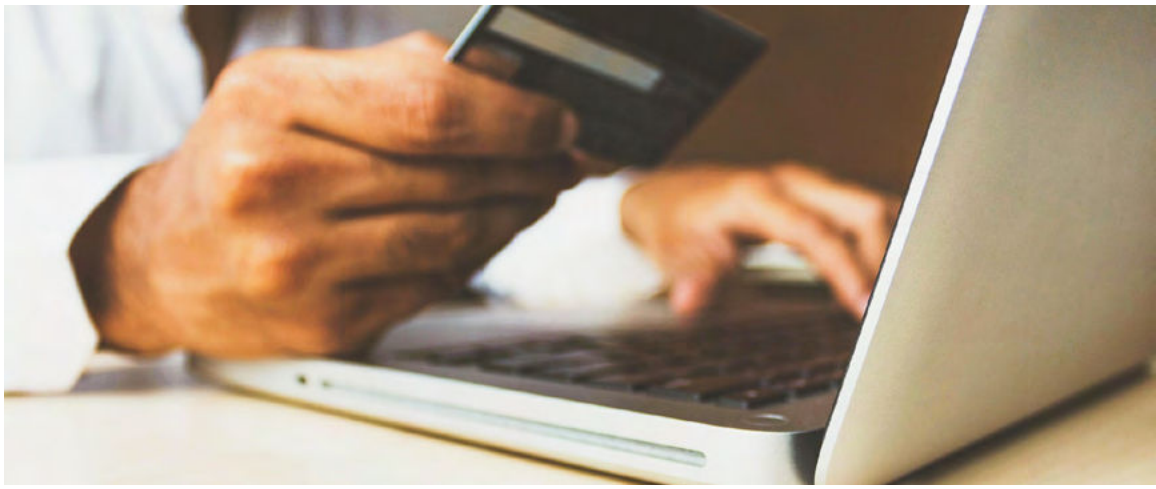

The "If X, Then Y" Quick Decision Guide

Use this simple framework before using any AI tool:

IF this is true...	THEN do this
The data contains names, SINs, health info, or can identify someone	**DO NOT enter into any public AI tool. Period.**
The output will be used for a decision affecting someone's rights, benefits, or employment	Require substantive human review. Document the human decision-maker.
The output will be shared externally (clients, public, partners)	Verify all facts independently. Have a qualified person approve before release.
You want to use a new AI tool not yet approved	Submit a request to IT/Privacy Officer. Do not use any new tool until approved.
The AI tool will process data outside Canada	A Privacy Impact Assessment is required before use if using personal data (Law 25 requirement).

ROLES AND RESPONSIBILITIES

A policy without clear ownership is a policy that gathers dust. For small organizations, roles may be combined. What matters is that each responsibility is assigned to someone.

ROLE	RESPONSIBILITIES	SMALL ORG ASSIGNMENT
Executive Leader	Approves policy; accountable to board; champions responsible AI use	ED, CEO, CAO… If a municipality, the mayor
Privacy Officer	Ensures privacy compliance; conducts PIAs; handles access requests. **REQUIRED under Law 25**	Default: highest authority. May delegate in writing
Policy Owner	Drafts and maintains policy; coordinates reviews; tracks updates	Operations Manager, IT lead, or designated staff
IT/Security Lead	Evaluates tools; manages approved list; monitors security	External IT provider or staff with tech responsibilities
All Staff	Follow policy; verify outputs; report concerns; complete training	Everyone including contractors and volunteers

Law 25 Requirement

You must publish the title and contact information of your Privacy Officer on your website.

CREATING YOUR POLICY: STEP-BY-STEP

Step 1: Assess Your Current Situation

- Survey staff about current AI tool use (formal or informal)
- Identify tools already in use and their purposes
- Review existing policies (privacy, IT, HR, communications)
- Consider the sensitivity of data your organization handles

Step 2: Define Your Guiding Principles

Your policy should reflect your organization's values. Consider adopting principles such as:

1. **Human Oversight:** AI assists but does not replace human judgment for consequential decisions.
2. **Privacy Protection:** Personal information is never entered into public AI tools.
3. **Accuracy Verification:** All AI outputs are verified before use.
4. **Transparency:** We are honest about when AI tools are used.
5. **Accountability:** Clear responsibility for AI-related decisions.
6. **Non-Discrimination:** AI use must not create or reinforce discrimination.

Step 3: Establish Clear Rules

Your policy must include specific rules in these categories:

Absolute Prohibitions (Never Permitted)

- Entering personal information into public AI tools
- Using AI for final decisions about individuals without human review
- Sharing confidential organizational information with unapproved AI tools
- Publishing AI-generated content without human review and approval
- Using AI for biometric analysis without CAI notification and consent
- Using personal accounts for organizational work with AI tools
- Installing AI browser extensions without IT approval

Conditional Permissions (Allowed with Safeguards)

- Using AI for drafting — with subject matter expert review before finalization
- Using approved enterprise AI tools — for tasks involving non-sensitive data
- AI-assisted decision support — with documented human review and final decision

General Permissions (Allowed)

- Learning about AI through training and experimentation with non-sensitive content
- Brainstorming, ideation, and early-stage research
- Editing and proofreading non-confidential text

Step 4: Define Approval Processes

RISK LEVEL	EXAMPLES	APPROVAL	REQUIREMENTS
Low	Internal drafts, research, brainstorming	Self-approval within policy	Follow policy; verify outputs
Medium	External communications, new tool requests	Manager or Policy Owner	Document use; expert review
High	Decision support, new AI systems, personal data	Executive + Privacy Officer	PIA required; human-in-the-loop

Step 5: Training and Awareness

A policy only works if people understand it. Include:

- Initial training when policy is launched
- Onboarding training for new staff
- Annual refreshers and updates when policy changes

Step 6: Review and Update Schedule

Generative AI evolves rapidly. Your policy needs:

- Formal review at least every 12 months (more frequently in the first year)
- Monitoring of regulatory developments from CAI and privacy commissioners
- Version number and effective date on the policy document

SAMPLE POLICY STATEMENTS AND RULES

Adapt these sample statements for your organization's policy. Items in [brackets] should be customized. Feel free to improve them.

Purpose Statement

> *"This policy establishes guidelines for the responsible use of generative artificial intelligence (AI) tools at [Organization Name]. It aims to capture the productivity benefits of these tools while protecting privacy, ensuring accuracy, maintaining legal compliance, and preserving public trust."*

Scope Statement

> *"This policy applies to all employees, contractors, volunteers, and board members when conducting work on behalf of [Organization Name]. It covers all generative AI tools, whether provided by the organization or accessed personally, when used for organizational purposes or with organizational data."*

Personal Information Prohibition

> *"Personal information must never be entered into public generative AI tools. This includes names, contact information, case details, health information, financial information, or any other information that could identify an individual. This prohibition applies regardless of whether the individual is a client, member, employee, or member of the public. Violation of this rule may result in disciplinary action."*

Verification and Human Review

> *"All content generated by AI tools must be reviewed by a qualified staff member before being used in any official capacity. The staff member using the tool is responsible for verifying the accuracy of all outputs. AI-generated content must not be used as the sole basis for decisions affecting individuals. The claim that 'the AI made an error' is not an acceptable justification for inaccuracies or harms."*

Automated Decision-Making (Law 25 Compliance)

> *"Any use of AI to assist in decisions that affect the rights or interests of individuals (e.g., hiring, service eligibility, performance evaluation) must include meaningful human review. The reviewing person must have the authority and competence to disagree with the AI suggestion. Individuals have the right to know when automated processing was used and to present observations to a human decision-maker. This requirement applies in accordance with Law 25, Section 12.1."*

Approved Tools

> *"The following AI tools have been approved for organizational use: [list tools]. Use of other AI tools for organizational purposes requires written approval from [designated role]. The approved tools list will be reviewed quarterly. Using personal accounts or unapproved free versions of AI tools for organizational work is prohibited."*

Transparency Statement

> *"When asked, [Organization Name] will be transparent about its use of AI tools. AI-generated content that interacts directly with clients or the public (e.g., chatbots) must be clearly identified as automated at the start of the interaction."*

INCIDENT RESPONSE AND REPORTING

Your policy should include procedures for handling AI-related incidents.

Types of AI-Related Incidents

- **Data Exposure:** Confidential or personal information entered into public AI tools
- **Accuracy Failures:** AI-generated errors that were published or acted upon
- **Bias Incidents:** AI outputs that resulted in discriminatory treatment
- **Security Breaches:** Unauthorized access, prompt injection, or other security issues
- **Policy Violations:** Use of unapproved tools or prohibited uses

Reporting Procedure

1. **Immediately stop** the activity that caused or discovered the incident
2. **Document** what happened, when, what data was involved, and what actions were taken
3. **Report** to your supervisor and the Privacy Officer within [X hours]
4. **Preserve** any evidence (screenshots, logs, records)

Law 25 Breach Notification

If a confidentiality incident involving personal information presents a risk of serious injury, you must notify the CAI and affected individuals. Keep a register of all confidentiality incidents.

QUICK-START IMPLEMENTATION

Don't let perfect be the enemy of good. Start with basics and build from there.

If You Have One Day

1. Send a brief survey asking what AI tools staff currently use
2. Communicate one rule: "Never enter personal information into public AI tools"
3. Identify who will be responsible for developing a full policy

If You Have One Week

1. Complete the self-assessment checklist in Section 10
2. Draft a basic policy using the sample statements in Section 7
3. Review them with your Privacy Officer
4. Get leadership approval and communicate to all staff

If You Have One Month

1. Conduct full AI tool inventory across the organization
2. Review existing policies for gaps and conflicts
3. Develop comprehensive policy with stakeholder input
4. Create approved tools list with vendor due diligence
5. Develop training materials and schedule rollout
6. Establish incident reporting procedures

Ongoing maintenance

1. Schedule regular policy review (AI tools and regulations change a lot!)
2. Monitor updates from the CAI and privacy commissioners
3. Track incidents and adjust policy based on lessons learned
4. Include AI policy in new staff onboarding

About Open North

Building trust in data, for the common good

Open North is a not-for-profit organization dedicated to advancing the common good. Working alongside governments and civic-minded organizations of all sizes, we provide data expertise to enhance decision making, drive innovation, improve public and civic services, and address society's most pressing challenges.

Our work focuses on building the capacity of organizations to make better decisions about managing their data so that it is useful, actionable, secure, and trustworthy throughout its entire lifecycle. At Open North, we combine deep expertise in data with a multidisciplinary approach. Our team includes urban planners, software engineers, community organizers, data scientists, cybersecurity and IT audit specialists, sociologists, geographers, and technology lawyers, bringing diverse perspectives to every project.

Open North is part of Montréal in Common, a project led by the City of Montréal as part of the Smart Cities Challenge, carried out with the financial support of the Government of Canada.

opennorth.ca

About the Smart Cities Challenge and Montréal in Common

Montréal in Common is an innovation community led by the City of Montréal whose partners are experimenting with solutions regarding access to food, mobility, and municipal bylaws, with a view to rethink the city. Montréal in Common projects are made possible thanks to the prize awarded to the City of Montréal by the Government of Canada as part of the Smart Cities Challenge.

Author: Cristiano Therrien

Canada Montréal SMART CITIES CHALLENGE Montréal en commun

ANNEXES: CHECKLISTS AND REFERENCES

Annex A - Self-Assessment Checklist

Use this checklist to assess your organization's readiness. Check each item you can answer "yes" to.

Governance and Accountability

- [ ] We have designated someone responsible for AI governance
- [ ] We have a designated Privacy Officer (required under Law 25)
- [ ] Leadership has discussed generative AI use
- [ ] We know which staff members are currently using AI tools

Policy and Documentation

- [ ] We have a written policy or guidelines on generative AI use
- [ ] The policy clearly states what is and is not permitted
- [ ] The policy addresses personal information in AI tools
- [ ] The policy requires human review of AI outputs
- [ ] The policy has a review date and version number

Privacy and Legal Compliance

- [ ] We have reviewed our obligations under Law 25/PIPEDA
- [ ] We prohibit personal information in public AI tools
- [ ] We have a process for Privacy Impact Assessments
- [ ] We know where our AI tools store and process data

Security and Access

- [ ] We have an inventory of AI tools being used
- [ ] We control access through organizational accounts
- [ ] We have reviewed terms of service for AI tools we use
- [ ] We have incident reporting procedures

Quality and Verification

- [ ] We require human review before external use
- [ ] We prohibit AI for decisions without human oversight
- [ ] Staff understand hallucination risks

Training

- [ ] Staff have received AI policy training
- [ ] Staff know who to contact with questions

Scoring

If you checked fewer than 10 items, develop a comprehensive policy. If 10-15, you have foundations but gaps to address. If 16 or more, focus on continuous improvement.

Annex B - Vendor Due Diligence Checklist

Before adopting a new AI tool, ask these questions and reflect on the answers:

1. Where is data stored and processed? Is it in Canada?
2. Will our data be used to train the AI model?
3. What security certifications does the vendor have (SOC 2, ISO 27001)?
4. How long is data retained? Can we delete it?
5. What are the breach notification obligations?
6. Does the vendor carry cyber liability insurance?

Check the **privacy and cybersecurity guides** from Open North for more information

References

Canadian regulatory guidance

- Commission d'accès à l'information du Québec (CAI)
- Office of the Privacy Commissioner of Canada - priv.gc.ca
- Government of Canada Guide on the Use of Generative AI
- Directive on Automated Decision-Making (Government of Canada)

International frameworks

- OECD AI Principles
- UNESCO Recommendation on the Ethics of AI
- ISO/IEC 42001:2023
- NIST AI Risk Management Framework

Open North resources

- Privacy Principles and Practices for Compliance with Law 25 - Open North guide on Quebec privacy requirements
- Cybersecurity for Small and Medium Organizations - Open North guide on cybersecurity fundamentals
- Data Privacy Policy: How to Write a Complete Draft - Open North guide on privacy policy development

Generative AI in Your Organization