

IN-DEPTH GUIDE

RISKS AND GOVERNANCE

Privacy and artificial intelligence: where your data goes



GXN

Senior expertise · Digital governance

Direct response

Limits, risks and decisions

July 30, 2026

Updated



ARTIFICIAL INTELLIGENCE · RISKS AND GOVERNANCE

Privacy and artificial intelligence: where your data goes

COMPREHENSIVE STUDY · GXN

GXN offers technical implementation and operational support. Legal interpretation should come from qualified legal counsel.*

TL;DR

The answer in one minute

- Your data leaves your company as soon as an employee sticks a document into a mainstream AI tool. It's already happening at your house.
- Offers for businesses generally exclude training on your data. Free personal accounts often don't. The difference is major.
- Five questions are enough to evaluate any supplier. They are lower.
- Canadian hosting is possible and relevant in many cases, without being a universal obligation.
- The biggest privacy risk in 2026 is not which provider you choose. This is the unregulated use that already exists.

01 · BACKGROUND GUIDE

The real portrait in Quebec companies

I start there because that's what I see in almost every mandate.

Management has not approved any AI tools. Officially, the company does not make any. In fact, several employees use consumer services on personal accounts, with company documents inside. A quote to be reformulated. A contract to summarize. A customer complaint to be answered politely.

Nobody's talking bad faith. They are effective people who try to do their job well. But each of these actions sends information from your company to a service you have not evaluated, under conditions that no one has read.

This is the real issue of confidentiality. Not the project you haven't started yet. The use which already exists and which no one regulates.

The journey of data, step by step

To judge a risk, you need to know what is really happening. Here is the path of a document you submit to a cloud AI system.

- **The departure.** The text leaves your station or server, encrypted in transit.
- **The treatment.** It arrives at the supplier, who cuts it and gives it to the model. The model produces a response.
- **Conservation.** According to the contract, the content can be kept for a few days for reasons of security and abuse, or not at all. On business offers, this duration is usually short and documented.
- **Training.** The critical point. Does your content serve to improve the model? On serious business offers, no, and it is written in the contract. On free personal accounts, the answer is often yes, or varies depending on the settings.
- **The return.** The answer comes back to you.

At no time does the template "hold" your document as a database would. But the information has indeed travelled, and it is this journey that you must frame.

The journey to be able to document



The five questions to ask any supplier

01

Use the same. They work with a model provider like an integrator.

1. Is my data used to train your models? The acceptable response is a written no to the contract, not a verbal assurance. If a box needs to be checked to get this no, make sure it is checked.

02

2. How long do you keep my content, and where. Look for a specific duration and a precise geographic location. “Our cloud partners” is not an answer.

3. Who in your country can access my content, and under what circumstances. There is almost always human access possible for security or support reasons. You want to know in which cases and under what controls.

03

4. What happens if I cancel. Deletion within which time, with which confirmation.

5. What contractual commitments do you agree to make in writing. This is the question that separates true enterprise providers from the rest. A general public service will not sign anything special for you. A corporate supplier, yes.

04

A company that gets five clear and written answers is in a strong position. One who gets five vague answers should look elsewhere.

04 · BACKGROUND GUIDE

The three possible architectures

Your level of control depends on the structure chosen. There are three, at least the most controlled.

General public service. An employee uses a public tool with an individual account. Minimum control, no visibility for the company, generic conditions. Acceptable for non-sensitive content. Unacceptable for everything else, yet it is the most common architecture at the moment, by default.

Company offer with contractual commitments. The company contracts directly, with non-training, limited retention and sometimes localization clauses. Reasonable control, moderate implementation effort. This is the appropriate solution for the majority of Quebec SMEs, including for sensitive business data.

Controlled infrastructure or hosted model. The processing takes place in an infrastructure that you control, possibly in Canada, with an open model. Maximum control. Real cost and complexity. This is justified for health information, employee files, information covered by a professional order, or industrial secrets.

Many companies end up with a combination of the latter two, which avoids paying the price of the most restrictive case everywhere.

THE LEVEL OF CONTROL ACCORDING TO THE ARCHITECTURE		
Architecture	Control	Appropriate context
General public service	Minimal	Non-sensitive content
Business offer	Reasonable	Framed business data
Controlled infrastructure	Maximum	Highly sensitive data

Law 25 and artificial intelligence

What this implies technically, in plain language, with the reminder that the legal interpretation belongs to your advisor.

When an AI system processes personal information, the same principles apply as for any other processing. In concrete terms, it raises technical questions that you need to be able to document.

- **What personal information actually enters the system.** Often more than expected, because an "anonymous" document contains a name in a signature or address in a footer.
- **Where they are treated and preserved.** Localization question, to be documented.
- **Who has access to it.** At your place, and at the supplier's.

- **Communication outside Quebec.** If treatment is done elsewhere, it triggers evaluation obligations. Your advisor will tell you which ones apply to your case.
- **The right of access and withdrawal.** If an individual requests the deletion of their information, you can do so in your AI system, including search indexes.
- **Minimization.** Please send more information than necessary to the model. Often, yes, for ease of design.

These are technical questions to which GXN answers through implementation: filtering before sending, logging, access control, flow documentation. The legal qualification of your obligations is up to your advisor.

06 · BACKGROUND GUIDE

Canadian Hosting: When it Matters

Honest answer: sometimes a lot, often not.

It really counts in four situations. When you process health information or data covered by a professional order. When your business customers or calls for tenders contractually require it. When you are a public body or work for one. When your sector has strong expectations, like financial services.

It means less to marketing content, public documentation or internal processes without personal information.

What matters in all cases is being able to answer the question. The day a business customer asks you where their data goes when you use AI, “I don’t know” is a bad answer. Documenting the journey, even when it passes elsewhere, is better than knowing nothing.

07 · BACKGROUND GUIDE

Ghost AI, your number one risk

Let's go back to the real problem.

You can choose the most rigorous supplier on the market, sign the best clauses, host in Canada. If three employees continue to contract in a public tool on their personal account, all this work is bypassed.

What works, in order.

- **Look before you ban.** Ask, without threat, who uses what and why. You'll get a free map of your internal frictions.

- **Provide an approved option.** Prohibition without solution produces clandestinity. A sanctioned tool, even imperfect, brings the use back to light.
- **Write a short policy.** Three pages that people read beat twenty pages that sleep. What is allowed, what is forbidden, what to do in case of doubt.
- **Train on concrete cases.** No theory. Examples from their actual work, with what is acceptable and what is not.
- **Review periodically.** The tools change, as do the uses.

What I see on the ground

The most common reflex after becoming aware is a total ban on emails. It lasts about six weeks. Then, usage resumes, more discreetly than before, and the company has lost its only chance to know what is going on.

The other mistake is to believe that the problem comes from the tool. It comes from the absence of decision. A company that has decided what categories of information can be released, and that has written it simply, manages its risk better than a company that has purchased the most secure solution without ever defining a rule.

When this guide is not enough

Three situations require proper accompaniment, with your legal counsel in the room.

- You process health information, employee data or information covered by a professional order.
- You are subject to specific sectoral or contractual requirements.
- You are about to sign a multi-year commitment with an AI provider.

GXN offers technical implementation and operational support. Legal interpretation must come from qualified legal counsel.

QUESTIONS AND ANSWERS

Frequently asked questions

01 Do our data train the models?

On business offers from large suppliers, usually not, and it is contractual. On personal and free accounts, it is variable and often yes by default. Check rather than assume.

02 Can be used with personal information.

Technically yes, with appropriate protections. The exact obligations that apply to your situation are the responsibility of your legal counsel. Technically, it involves minimizing, controlling access, logging and erasing capacity.

03 **It is absolutely necessary to stay in Canada.**

Not universally. This becomes important for certain categories of data and contractual requirements. What is universal is knowing and documenting where your data is going.

04 **How to prevent our employees from using unapproved tools.**

You will not prevent them from doing so alone. Provide an approved option, write a short policy, form on real cases, and review regularly.

05 **One model can reveal our data to another customer.**

With a company offer that excludes training, this risk is very low. It increases with consumer services and with poorly designed configurations, especially when several clients share the same poorly partitioned document index.

06 **What to ask an integrator before signing.**

The five questions above, plus a sixth: show me the diagram of the journey of my data, from the workstation to the response. If he can't draw it, he didn't design it.

EDITORIAL METHOD

Background guide prepared by GXN (Digital Governance), a division of MD79. Revision planned according to the rhythm documented for this cluster.

OFFICIAL SOURCES

References to be checked according to your situation

These references support the external rules and frameworks cited in this guide. They are not a substitute for legal or professional advice tailored to your organization.

NIST

Artificial Intelligence Risk Management Framework



Quebec Access to Information Commission

Major changes made by the Act 25



ABOUT THE PUBLISHER

GXN

A directly accessible senior expertise, with more than twenty years of experience in technology implantation and a master's degree in artificial intelligence within the company.

[Discover GXN](#) ↗

NEXT STEP

GXN offers technical implementation and operational support. Legal interpretation must come from qualified legal counsel.

You want to know where your data is actually going today.

The AI diagnosis includes flow mapping, usage already in place in the company, and the architecture appropriate to your level of sensitivity.

No call required. No sales sequence.