

GOVERNANCE • SMBS AND ORGANIZATIONS • AI

Practical *AI* governance guide for SMBs and organizations.

Structure usage, protect data, and identify the right automation opportunities.

This guide is designed for SMB and organizational leaders who have already downloaded and completed the **seven-question Starter Pack** and want to go further. The Starter Pack fits on a single page and covers 80 % of the need. This guide structures the remaining 20 %.

You can use it as a discussion outline with your IT team and your operations lead. You can also present it to your board of directors, adapting the language to your vocabulary and your reality. This is not a legal document. Your legal, cybersecurity, and privacy specialists remain the ones who validate final compliance.

VERSION

1.0 • May 2026

AUTHOR

Steve Johnston

FORMAT

Printable • Letter,
portrait

COMPANION

7-Question Starter
Pack

How to use this guide in 30, 60, or 90 minutes

HOW TO USE

Three paths based on the time you have. Choose the one that fits your week, not your ambition.

30 MIN PATH • DISCOVER

Find where to focus your energy.

Read the how-to, then sections **1 (Objective)**, **4 (Data classification)**, and **9 (Human validation)**.

You know where to put your energy this week.

60 MIN PATH • FRAME

Set the foundations and the framing.

All sections **1 through 7** (foundations, principles, classification, uses, registry, approval process).

You leave with a draft charter in hand.

90 MIN PATH • STRUCTURE

Cover the full framework.

The full guide, from the how-to through the practical appendices.

You have your complete framework, ready to validate with your specialists.

PRODUCTIVITY THESIS

Governing AI does not mean slowing innovation.

Good AI governance is not built to say no. It is built to say yes faster to the right use cases.

CONCRETE BENEFITS OF A CLEAR FRAMEWORK

- Avoid bad uses
- Accelerate good uses
- Reduce manual data entry
- Standardize processes
- Improve data quality
- Free up human time
- Prioritize profitable cases

AI maturity grid

WHERE YOU STAND

Before diving into the 14 sections, locate your organization. This grid points you to your next concrete step.

LEVEL	DESCRIPTION	NEXT STEP
0	No framework, scattered uses.	Complete the seven-question Starter Pack.
1	Basic rules in place.	Build the list of approved tools.
2	Active use-case registry.	Prioritize 3 high-impact cases.
3	Structured governance.	Train the teams.
4	AI integrated into processes.	Measure gains and risks.

Most SMBs I meet sit between levels 0 and 2. Moving from level 0 to level 1 is quick: that is precisely the purpose of the Starter Pack. Moving from level 1 to level 2 requires this guide. Beyond that, it is a matter of operational discipline, not added complexity.

01. Objective

SECTION 1 OF 14

The objective of this charter is to define clearly why your organization uses artificial intelligence, within what limits, and under what rules. This section answers a simple question: "Why are we doing this?"

SAMPLE WORDING TO ADAPT

AI is used in our organization to improve our teams' productivity, the quality of our decisions, our customers' experience, and the efficiency of our operations. It complements human judgment, never replaces it blindly. This charter sets the guardrails that protect our data, our reputation, our employees, and our customers as we draw value from these tools.

WHY THIS SECTION MATTERS

Without a clear objective, each team interprets AI's place differently. With a clear objective, later trade-offs become simpler.

02. Scope

SECTION 2 OF 14

Scope defines who and what this charter applies to.

WHO

- Full-time and part-time employees
- Consultants and contracted self-employed workers
- Executives and members of the leadership team
- Vendors who process our data or operate on our systems

WHAT

- Internal AI tools built in-house or integrated into our systems
- External AI tools used at work (ChatGPT, Copilot, Gemini, Claude, Perplexity, and equivalents)
- AI automations in sales, marketing, operations, HR, customer service, and finance
- Any function where AI influences a decision or a deliverable of the organization

USEFUL DISTINCTION BY USE TYPE

USE TYPE	EXAMPLE	CONTROL LEVEL
Personal assistance	Rewording an email, summarizing an internal note	Light
Operational	Analyzing customer data, generating a monthly report	Medium
Decision support	Helping reject an application, prioritizing customers, setting a price	High
External automation	Customer chatbot, automated recommendation sent to a third party	Very high

The higher the control level, the more rigorously sections 9 (human validation) and 10 (risk management) must apply.

03. Guiding principles

Six to eight principles are enough to anchor a culture of AI use in the organization. More than that and no one remembers them.

RECOMMENDED PRINCIPLES

1. **Real usefulness.** Every AI use must solve a concrete business problem. No AI for the sake of AI.
2. **Data protection.** No confidential, personal, or sensitive data should be entered into an unapproved AI tool.
3. **Accountable human.** A person remains accountable for the result, even when AI helped. This is the rule that keeps everything else from drifting.
4. **Reasonable transparency.** Customers, employees, or partners must be informed when AI influences an important interaction.
5. **Accuracy and validation.** AI outputs must be verified before being used in a decision, an official document, or an external communication.
6. **Fairness and non-discrimination.** AI must not create or amplify bias, especially in HR, credit, pricing, customer service, or prioritization.
7. **Security.** AI tools follow the organization's cybersecurity rules, validated with your IT lead.
8. **Traceability.** Important uses are documented: tool used, objective, data involved, owner, validation.

These principles align with the general expectations of Canadian authorities on privacy protection and responsible use of generative AI. Your specialists can tie them to the specific provisions of the law that applies to your organization.

04. Data classification

Probably the most useful section for an SMB or organization. It defines what can go into an AI tool and what cannot.

FIVE-LEVEL GRID

CATEGORY	TYPICAL EXAMPLES	PUBLIC AI ALLOWED?
Public	Website content, marketing brochure, published article	Yes, no restriction
Internal	Internal procedures, document templates, general strategic ideas	With care, approved tool only
Confidential	Financial statements, contracts, strategy, customer lists, negotiated prices	No, unless approved tool on a Business/Enterprise plan
Personal	Name, email, employee file, identifying customer data	No, unless a formal framework with prior assessment
Sensitive	Health, personal finances, sensitive HR data, minors	Highly restricted, formal framework required

WORKED EXAMPLE

Web visitor → form → Excel sheet → human → CRM

Our visitor form captures **name, email, and customer need**. This data falls into the **Personal** category. Direct consequence: no public AI tool can process it without a formal framework and a prior assessment. This finding shapes every decision in the following sections for this use case.

Adaptation work: rewrite these categories using the vocabulary and real examples from your organization. A foundation does not handle the same data as a professional services firm or a manufacturer. The grid only works if your teams recognize their own data inside it.

05. Permitted, supervised, and prohibited uses

Be very concrete. The grid below is a starting point to personalize.

GENERALLY PERMITTED USES (NO APPROVAL NEEDED, FOR ALL EMPLOYEES)

- Rewording an internal text or an email
- Generating ideas for a presentation, a campaign, or a process
- Creating drafts to be validated afterward
- Summarizing a non-confidential document
- Producing project plans or analysis structures
- Helping structure a line of thought, a question, or a request

SUPERVISED USES (PRIOR VALIDATION BY A DESIGNATED OWNER)

- Analyzing customer or employee data
- Producing a strategic recommendation to leadership
- Generating legal, financial, or HR content, even in draft form
- Automating a customer response (chatbot, automatic email)
- Connecting AI to a CRM, an ERP, or a shared Excel file
- Automated scoring or automated sorting of prospects, candidates, or customers

PROHIBITED OR STRONGLY RESTRICTED USES

- Entering personal or sensitive data into an unapproved public AI tool
- Letting AI make a final decision about an employee, a customer, or a vendor without human review
- Generating false references, testimonials, or results
- Using copyrighted content without clearance
- Automating an important decision without documented human oversight
- Using AI to monitor employees without a formal framework and consent
- Circumventing the organization's existing security policies

WORKED EXAMPLE

Web visitor → form → Excel sheet → human → CRM

Goal for this case: **reduce manual data entry** between the Excel sheet and the CRM. Because the data falls into the Personal category (see section 4), this case sits in the **supervised** zone: prior validation by a designated owner, not a free-for-all use.

This section is the one your employees will read first. The clearer it is, the more it will be respected.

A living table, not a frozen document. One line per important use case in the organization. Maintained by a named owner (see section 8) and reviewed quarterly.

RECOMMENDED COLUMNS

FIELD	EXAMPLE
Use-case name	Automated summary of contact-form submissions from the website
Department involved	Sales
Person accountable	Director of Sales
Tool used	ChatGPT Business or Microsoft 365 Copilot
Data used	Web form, prospect name and email
Personal data?	Yes
Risk level	Medium
Human validation in place?	Yes, the director validates each summary before transfer to the CRM
Expected benefit	4 hours per week saved on qualification
Status	Approved (March 2026)

WORKED EXAMPLE

Web visitor → form → Excel sheet → human → CRM

Here is the actual line for this case in the registry:

FIELD	VALUE
Use-case name	Assisted transfer of visitor form to CRM
Department involved	Sales / Operations
Person accountable	Operations lead
Tool used	Approved automation platform (Business plan)
Data used	Name, email, customer need (Personal)
Personal data?	Yes
Risk level	Medium
Human validation in place?	Yes, on exceptions (ambiguous forms)
Expected benefit	Less manual data entry, fewer errors, faster follow-up
Status	Pilot · May 2026

This registry becomes the operational inventory of AI in your organization. It also quickly reveals the cases that should have gone through the approval process and did not. It is a useful document to present to an auditor, an institutional client, or a partner asking how you manage your data.

Before adopting a new AI tool in the organization, the designated owner checks several points.

APPROVAL CRITERIA

- Where is the data hosted geographically?
- Is the data used to train the vendor's models?
- What are the terms of service (read them in full, not just the summary)?
- Which privacy options are available?
- Does the tool integrate with our existing systems?
- What are the access rights and who controls them?
- Are activity logs available?
- Can we delete our data on request?
- Does the tool meet our compliance obligations?
- What is the total cost: licences, integration, training, support?

SIMPLE RULE TO ENGRAVE IN THE CHARTER

No new AI tool may be used with customer, employee, or financial data without prior validation by the designated owner.

Even in an SMB, you have to name people. Often, several roles fall to the same person. That is not a problem, as long as the responsibilities are named explicitly.

TYPICAL ROLES

ROLE	PRIMARY RESPONSIBILITY
Leadership	Approves the AI direction and the acceptable risks
AI lead	Maintains the registry, oversees use cases, monitors approved tools
Privacy lead	Validates the risks tied to personal data
IT / cybersecurity lead	Validates access, integrations, and security
Department managers	Oversee day-to-day uses within their team
Employees	Use AI according to the rules and validate the outputs
Vendors	Meet the organization's AI and data requirements

Your privacy governance obligations already exist within your organization. AI governance hooks into them, rather than living in a parallel document.

This is the ethical and legal anchor of the guide. The principle: AI can recommend, structure, summarize, or automate part of the work, but final accountability rests with a designated human.

HUMAN VALIDATION REQUIRED FOR

- Decisions affecting a customer (acceptance, refusal, terms, special pricing)
- HR decisions (candidate selection, evaluations, disciplinary measures, terminations)
- Public communications (official notices, press releases, public positions)
- Contracts, offers, and bids sent externally
- Financial data (results, forecasts, regulatory filings)
- Strategic recommendations presented to the board or leadership
- Automations connected to critical systems (CRM, ERP, accounting)
- Any content generated from sensitive data

USEFUL WORDING TO CARRY INTO THE CHARTER

*AI can recommend, structure, summarize, or automate part of the work. For the decisions listed above, the final decision rests with **the designated human**. That person is accountable for the result and its consequences.*

WORKED EXAMPLE

Web visitor → form → Excel sheet → human → CRM

In our case: **the human validates the exceptions** (ambiguous form, non-standard request, incomplete data). The normal flow passes through without intervention. That is precisely what makes the automation acceptable from a governance standpoint: a human remains accountable for cases that fall outside the framework.

This wording matters in practice: it keeps you outside the regime of decisions "based exclusively on automated processing" contemplated by certain privacy laws, as long as a human truly intervenes in the final decision.

A simple grid is enough for most organizations.

RISK	CONCRETE EXAMPLE	CONTROL MEASURE
Confidentiality	Customer data entered into a free-tier AI tool	Approved tools only, Business or Enterprise plan
Hallucination	Fabricated content in a proposal sent to a customer	Human validation required before external delivery
Bias	Automated candidate sorting that replicates a historical bias	HR review required before any decision
Security	AI connection to the CRM or ERP without controls	Access control, activity logs turned on
Vendor dependency	A critical process resting on a single external tool	Documented contingency plan, alternative identified
Reputation	Incorrect automated response sent to a customer	Prior testing, supervision during the pilot phase
Legal	Faulty contract advice generated by AI and used as is	Validation by a qualified professional required

WORKED EXAMPLE

Web visitor → form → Excel sheet → human → CRM

Risk identified: personal data flowing across several systems, with a chance of data-entry error or mis-routing. **Control applied:** human validation on exceptions, activity logs turned on within the automation platform, quarterly review of the registry.

This grid is not exhaustive. Your specialists in cybersecurity, insurance, and risk can enrich it based on your organizational profile.

11. Transparency toward customers and employees

SECTION 11 OF 14

The more AI influences an important decision, the stronger transparency must be.

THREE WORDINGS TO ADAPT TO THE CONTEXT

- *"Some responses may be assisted by an artificial intelligence tool."*
- *"Your data may be processed with digital tools to improve our service."*
- *"Any important decision remains reviewed by an authorized person."*

THREE KEY MOMENTS WHERE TRANSPARENCY IS REQUIRED

1. At initial data collection (privacy policy, contact form)
2. During an AI-assisted interaction (a customer chatbot, for example, that identifies itself as an automated tool)
3. When a decision is communicated to the person concerned

This transparency is not only a legal obligation, it is also a mark of respect that builds trust over time.

A meaningful share of AI risk for an SMB comes from vendors. Before adopting an AI tool or an automated service provided by a third party, ask these nine questions and keep the answers in writing.

QUESTIONS TO ASK ANY AI VENDOR

1. Do you use our data to train your models? If so, can we opt out contractually?
2. Where is the data stored geographically?
3. Who can access the data on your side?
4. Do you offer an accessible activity log?
5. Can we delete our data on request?
6. Do you use subcontractors who handle our data?
7. Does the data leave Canada? If so, how is that transfer governed?
8. What security measures are in place (encryption at rest, in transit, access controls)?
9. What happens in the event of an incident, and how are you required to notify us?

WORKED EXAMPLE

Web visitor → form → Excel sheet → human → CRM

For our case: **vet the automation tool used to move the form to the CRM**. The nine questions apply as is. Keep the answers in your compliance file. Note in particular the data location and the option to disable model training on your content.

Keep the answers in your compliance file. They are part of the documentation you will present to an auditor, an institutional client, or a regulator should the question come up.

13. Minimum employee training

A charter is only useful when the team knows how to use it. A short training, at hire and then annually, is enough.

MINIMUM TRAINING CONTENT

- What generative AI is and its main limits
- What you can do in our organization (permitted uses)
- What you must never do (prohibited uses)
- How to write a good prompt to an AI tool
- How to validate an output before using it
- How to protect data during an AI interaction
- How to report a problem or an incident
- How to propose a new use case

The training can fit into one hour for a new employee. A fifteen-minute refresher each year is enough for the team in place, barring a major change in tools or rules.

14. Monitoring, audit, and continuous improvement

An AI charter must evolve. AI shifts every three to six months, your usage shifts with it.

RECOMMENDED CADENCE FOR AN SMB OR ORGANIZATION

ITEM	FREQUENCY
Review of the use-case registry	Quarterly
Review of approved tools (still appropriate?)	Twice a year
Review of incidents and lessons learned	After each incident
Charter update	Annual or after a major change
Employee training	At hire, then annually
Internal audit (AI lead)	Annual, formal
External audit (if applicable by sector)	Per sector-specific obligations

Document each review with a date and an owner. This is not bureaucracy: it is what proves, should a question arise, that your organization takes AI governance seriously.

Recommended practical appendices

CHARTER COMPANION

Beyond the 14 sections, your charter benefits from a few practical tools. These appendices live next to the charter, not inside it. They evolve faster than the principles.

- List of approved AI tools (kept current)
- List of prohibited or unapproved tools (with reason)
- New AI use-case request form
- Detailed data classification grid for your sector
- Risk assessment grid for your organizational profile
- Ready-to-fill AI registry template
- Checklist of questions to ask AI vendors
- Human validation policy (excerpt from section 9)
- AI incident procedure
- Sample security guidelines for your typical use cases

WORKED EXAMPLE • SUMMARY

Web visitor → form → Excel sheet → human → CRM

Measured gain after implementation: fewer data-entry errors, less delay between the form arriving and the entry into the CRM, better sales follow-up. All within a framework where the human stays accountable for exceptions and personal data stays protected. This is exactly what good AI governance aims for: saying yes faster to the right use cases.

How this guide connects to the Starter Pack

The seven-question Starter Pack covers 80 % of the need for most organizations. It fits on one page, takes an hour to draft, and gives your team a clear answer on what is allowed and what is not.

This practical guide structures the remaining 20 %: it serves when your organization grows, when complexity rises, or when your cybersecurity and privacy specialists want a more formal framework for their work.

The ideal way to use them: start with the Starter Pack this week. Put it in place. Six months later, when you have watched how your team actually uses AI, open this guide and adapt the sections that speak to you. Not all at once. One or two per month is plenty.

My role in all of this

I work in digital execution architecture and operational transformation. In that practice, **cybersecurity, data protection, and compliance run across everything**: they are an integral part of every architecture I design and every engagement I lead. This is not a domain external to my work, it is a built-in reflex, both

necessary and protective, embedded in the very design of my solutions.

I collaborate with sharp specialists in cybersecurity, privacy, and law for final compliance, because those domains demand specific expertise they master better than I do. The architectural framework I set up front already incorporates these concerns, which makes their work faster, more precise, and better anchored to your operational reality.

This guide is an accessible starting point. It does not eliminate the need for your specialists. It gives them a solid framework to support you better and faster.

TO MOVE INTO ACTION

Your *3 next actions* this week.

- 01 Identify your **5 current AI uses** in the organization (even informal ones).
- 02 Classify the **data used** in each one against the five-level grid (section 4).
- 03 Pick one **simple case to optimize** in the next 30 days, following the worked-example logic of this guide.

If you would like to talk it through for your organization, I'd be glad to.

CONTACT ME →

STEVEJOHNSTON.CO/CONTACT/ · CONFIDENTIAL DISCUSSION · NO COMMITMENT

Practical AI Governance Guide for SMBs and Organizations · Version 1.0 · May 2026

Steve Johnston · The Architect of Digital Execution Systems · stevejohnston.co

Companion to the seven-question Starter Pack · stevejohnston.co/starter-pack-gouvernance-ia/