



AI Risk Conversation Guide

A structured discussion for identifying business consequence, information exposure, human authority, and response before an AI workflow grows.

STRUCTURE - CLARITY - BETTER DECISIONS

Practical architecture for accountable AI-enabled work.

What this template is for

Use this guide to run a focused 45- to 60-minute conversation with leadership, staff, advisors, or a technology partner. The goal is a shared risk picture and a clear next decision, not a generic list of fears.

BEST USED WHEN

Proposed or existing AI uses that touch customers, employees, financial information, professional services, operational decisions, or external communications.

WHO SHOULD BE IN THE ROOM

Include the workflow owner, someone who performs the work, a subject-matter reviewer, and the person who can accept, reduce, transfer, or reject the risk.

CANADIAN CONTEXT

Confirm applicable federal, provincial, territorial, sector, professional, contractual, accessibility, and bilingual-service requirements. Use this workbook as operational preparation, not legal advice.

A five-step working rhythm

01

Business consequence

Describe who benefits, who may be affected, and what can be harmed if the workflow is wrong or unavailable.

02

Information exposure

Identify personal, confidential, privileged, financial, employee, and third-party information across the lifecycle.

03

Reliance + authority

State how people may rely on the output and which decisions or actions require a person.

04

Quality + inclusion

Test accuracy, completeness, consistency, accessibility, language, regional context, and unusual cases.

05

Evidence + response

Define monitoring, incident ownership, containment, correction, communication, and learning.

Map one risk conversation

Record the material decisions and open questions; do not try to capture every comment.

PROMPT	YOUR WORKING NOTES
01 Workflow, expected value, owner, and affected groups	
02 Most consequential failure and who could be harmed	
03 Sensitive information, access, provider use, and retention	
04 Human authority, reliance risk, challenge, and escalation	
05 Representative cases, limitations, and evidence threshold	
06 Incident response, correction, communication, and fallback	
07 Decision, risk owner, next action, and review date	

Separate insight from financial authority

STARTING FRICTION

An owner wants earlier cash-flow and margin signals, but incomplete feeds or confident summaries could distort an important decision.

WORKFLOW MAP



AI-SUPPORTED ROLE

Detect changes, summarize drivers, and pose questions using reconciled, permitted business data.

HUMAN AUTHORITY

The owner or accountant verifies figures and retains authority for payments, filings, financing, pricing, staffing, and commitments.

PILOT MEASURE

Unsupported claims, reconciliation differences, decision lead time, overrides, and material exceptions missed.

DISCUSSION PROMPT

What evidence would the reviewer need to trust, correct, reject, or escalate this output?

Assess risk in an AI-assisted matter summary

STARTING FRICTION

A concise summary is useful, but omission, stale law, incorrect chronology, or unintended disclosure can create serious consequences.

WORKFLOW MAP



AI-SUPPORTED ROLE

Organize facts and source references within the approved file; identify gaps and inconsistent dates.

HUMAN AUTHORITY

A lawyer validates privilege, relevance, legal interpretation, deadlines, advice, and every use outside the internal working record.

PILOT MEASURE

Material omissions, unsupported statements, review time, confidentiality incidents, and corrections after approval.

DISCUSSION PROMPT

What evidence would the reviewer need to trust, correct, reject, or escalate this output?

Route exceptions without hiding uncertainty

STARTING FRICTION

A risk or operations manager needs signals across incidents, vendors, controls, and follow-up, but aggregation can create false confidence.

WORKFLOW MAP



AI-SUPPORTED ROLE

Group related signals, surface anomalies, and prepare an evidence-linked brief with uncertainty visible.

HUMAN AUTHORITY

The manager decides severity, owner, response, communication, and whether the evidence is sufficient to close the issue.

PILOT MEASURE

Time to triage, false positives, missed events, reopened issues, and actions without adequate evidence.

DISCUSSION PROMPT

What evidence would the reviewer need to trust, correct, reject, or escalate this output?

Before the pilot grows

Use these prompts to test whether the proposed workflow can operate responsibly in its Canadian business context.

01

Discuss consequence, not novelty

Risk increases with the effect of the decision, the sensitivity of information, and the degree of reliance - not simply because AI is involved.

02

Map the full information path

Include prompts, retrieved sources, generated outputs, logs, vendor access, exports, backups, and deletion.

03

Make uncertainty operational

Define when the system must abstain, ask for missing information, route to a person, or stop.

04

Consider affected people

Test whether errors or exclusions may affect employees, customers, applicants, vulnerable groups, or people needing accessible or bilingual service.

05

Prepare the response before launch

Name the incident owner, containment step, correction path, notification decision, fallback, and evidence to preserve.

06

Revisit the risk

Review after incidents, near misses, provider changes, new data, expanded scope, or changes in Canadian legal and regulatory expectations.

CANADIAN REFERENCE POINTS

Office of the Privacy Commissioner of Canada - Privacy and artificial intelligence
priv.gc.ca/en/privacy-topics/technology/artificial-intelligence/

Government of Canada - Artificial Intelligence guidance, principles, and policy
canada.ca/en/services/science/innovation/artificial-intelligence.html

General operational guidance only. Confirm applicable legal, privacy, cybersecurity, professional, contractual, and regulatory requirements with qualified advisors.