



AI Incident Response Playbook

A worksheet for responding when AI gives harmful advice, exposes data, misroutes work, or fails in production.

STRUCTURE - CLARITY - BETTER DECISIONS

Practical architecture for accountable AI-enabled work.

What this template is for

Define severity, containment, ownership, communication, fallback, evidence capture, and post-incident improvement for AI-supported operations. Use this guide to make the operating choice explicit, test it with evidence, and record who owns the next decision.

BEST USED WHEN

Production AI systems and customer-facing assistants.

WHO SHOULD BE IN THE ROOM

Include the workflow owner, a person who performs the work, a reviewer with authority to approve or stop it, and the relevant data or tool owner.

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

Detect

Which signal shows the workflow is unsafe, degraded, or unavailable?

02

Contain

What must be paused, blocked, revoked, or routed to humans?

03

Communicate

Who needs to know internally or externally?

04

Recover

What fallback path restores operational continuity?

05

Learn

Which prompt, tool, policy, source, or review gate must change?

Build the working record

Complete one record for one workflow. Write enough to support a decision, not a generic technology wish list.


PROMPT	YOUR WORKING NOTES
01 Workflow affected	
02 Incident type	
03 Response owner	
04 Detect	
05 Contain	
06 Communicate	
07 Decision, owner, evidence, and review date	

Contain a cross-client disclosure

STARTING FRICTION

A generated brief includes information from a different engagement.

WORKFLOW MAP



AI-SUPPORTED ROLE

Surface the incident signal, affected outputs, source evidence, and a bounded containment checklist. Flag uncertainty around detect.

HUMAN AUTHORITY

The incident owner stops the workflow, assesses impact, communicates as required, and authorizes restart. Resolve the detect question for this case.

PILOT MEASURE

Track time to detect, contain, correct, notify where required, and prevent recurrence. Record detect exceptions separately.

DISCUSSION PROMPT


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

Correct a wrong customer reply

STARTING FRICTION

An assistant gives outdated refund information that a customer may rely on.

WORKFLOW MAP



AI-SUPPORTED ROLE

Surface the incident signal, affected outputs, source evidence, and a bounded containment checklist. Flag uncertainty around contain.

HUMAN AUTHORITY

The incident owner stops the workflow, assesses impact, communicates as required, and authorizes restart. Resolve the contain question for this case.

PILOT MEASURE

Track time to detect, contain, correct, notify where required, and prevent recurrence. Record contain exceptions separately.

DISCUSSION PROMPT


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

Respond to an unsupported figure

STARTING FRICTION

A draft summary cites a number that cannot be traced to the approved ledger.

WORKFLOW MAP



AI-SUPPORTED ROLE

Surface the incident signal, affected outputs, source evidence, and a bounded containment checklist. Flag uncertainty around communicate.

HUMAN AUTHORITY

The incident owner stops the workflow, assesses impact, communicates as required, and authorizes restart. Resolve the communicate question for this case.

PILOT MEASURE

Track time to detect, contain, correct, notify where required, and prevent recurrence. Record communicate exceptions separately.

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

Detect

Severity levels are defined. Name the owner and the operating decision before selecting a tool.

02

Contain

Containment steps can be executed quickly. Record the approved sources, permissions, and missing information.

03

Communicate

Fallback workflow is available. Test this boundary with representative and adverse cases.

04

Recover

Evidence can be preserved. Preserve review, escalation, and correction evidence.

05

Learn

Post-incident changes update the operating system. Set a measurable stop condition and the next review date.

06

Readiness to proceed

AI incident response is part of operating architecture. A team should know how to pause, explain, recover, and improve before failure happens.

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.