



TECHNOLOGY & OPERATIONAL INSIGHTS • AI & AUTOMATION

AI Readiness: Seven Questions Every Organization Should Answer Before Implementation

A boardroom-level readiness test covering use cases, data, oversight, security, evaluation, ownership, and ongoing operations.

READING TIME 4 min **CATEGORY** AI & Automation **PUBLISHED BY** Client Care Operations

PUBLISHED 20 October 2026

EXECUTIVE SUMMARY

- 01** The business question is not whether AI is powerful, but whether a specific system is useful, appropriate, controllable, and worth maintaining in your environment.
- 02** Leadership should be able to answer seven questions in plain language: scope, information, accountability, testing, security, operations, and the business case.
- 03** A readiness discussion should end in one of three legitimate decisions: a limited pilot, prerequisite data or process work, or a simpler non-AI solution. Privacy and security review belong in that discussion, not after it.

IN THIS ARTICLE

- 01** 1. What exact decision or activity are we improving?
- 02** 2. Is the underlying information suitable?
- 03** 3. Who remains accountable for decisions?
- 04** 4. How will we test reliability?
- 05** 5. What security risks does the design create?
- 06** 6. Who will operate and review the system after launch?
- 07** 7. What makes the business case worthwhile?
- 08** Turn the answers into a small, controlled pilot

Artificial intelligence can support a range of organizational tasks, from searching internal documents to classifying requests and drafting summaries. But the relevant business question is not whether AI is powerful. It is whether a specific system is useful, appropriate, controllable, and worth maintaining in a particular environment.

Before commissioning an AI integration, leadership should be able to answer seven questions in plain language.



Figure 1 Seven readiness questions lead to one of three decisions

Original CCO diagram based on this article.

If the team cannot define success, it cannot credibly claim improvement.

KEY INSIGHT

1. What exact decision or activity are we improving?

Start with one real workflow. Describe the input, desired output, responsible person, and what a successful outcome looks like. *Use AI to make operations smarter* is not a project scope. *Help managers find relevant policy sections when answering routine internal questions* is much closer.

Ask whether a simpler rule, improved search function, form, or conventional software integration could solve the same problem. NIST's AI Risk Management Framework encourages understanding the system's intended context and determining whether AI is appropriate for the use case.^[1]

2. Is the underlying information suitable?

Outputs depend on the completeness, relevance, accuracy, and permissions attached to the information they use. If documents contain conflicting policies or obsolete procedures, AI may reproduce that inconsistency in polished language.

Identify data owners, sensitivity, lawful uses, retention requirements, and the source documents that users should trust. Document what information may not be sent to third-party services. Canadian privacy obligations vary by organization, sector, province, and activity; the federal Privacy Commissioner provides general business guidance, but project-specific legal review may be necessary.^[2] Québec's private-sector privacy law, for example, requires enterprises subject to it to conduct a privacy impact assessment for projects to acquire, develop, or overhaul an information system involving personal information, and requires organizations to inform people when a decision about them is based exclusively on automated processing.^[4]

3. Who remains accountable for decisions?

A person should be able to identify when a system is wrong, correct its output, and stop or override the workflow. Set clear boundaries: can AI suggest, classify, summarize, or trigger an action? Can it modify a record? Can it send messages externally? The more consequential the action, the stronger the human approval and audit controls should be.

4. How will we test reliability?

Define representative real-world cases, known difficult cases, and what constitutes an unacceptable error. Measure more than whether an answer appears fluent. Depending on the use case, relevant checks include factual accuracy, unsupported claims, completeness, correct retrieval of authorized information, and whether escalation occurs when confidence is insufficient.

A pilot should include both routine and exceptional cases. If the team cannot define success, it cannot credibly claim improvement.

5. What security risks does the design create?

AI systems may introduce novel failure paths on top of ordinary application-security risks. OWASP's Top 10 for LLM and generative AI applications identifies issues including prompt injection, sensitive-information disclosure, excessive agency, and improper output handling; its 2026 edition updates the rankings and expands the threat coverage of the 2025 list.^[3]

Controls may include limiting tool permissions, validating outputs, protecting secrets, segregating data, recording actions, and restricting system authority. These protections must be implemented and tested for the proposed design; they are not properties of AI technology in general.

6. Who will operate and review the system after launch?

Identify the accountable owner for changes, model or vendor updates, incident response, feedback, monitoring, and policy revisions. If the workflow changes, who revises prompts, rules, evaluation cases, or connectors? What is the process for disabling the integration when needed?

A prototype is not an operational service until ownership, support, and monitoring are defined.

7. What makes the business case worthwhile?

Set a baseline and evaluate benefits alongside implementation effort, data preparation, exception-handling cost, supervision, risk, and vendor fees. Possible measures include turnaround time, manual touches, percentage of outputs accepted without correction, escalation rate, and user satisfaction.

Use conservative estimates. An AI-generated summary that saves two minutes but requires extensive review may not produce meaningful net value.

Turn the answers into a small, controlled pilot

A readiness discussion should result in one of three decisions: proceed to a limited pilot, complete prerequisite data/process work, or defer AI in favour of a simpler solution. Each is a legitimate outcome.

CCO PERSPECTIVE

Responsible AI integration requires a process, a person accountable for it, and evidence that the new approach is better than the old one.

Client Care Operations · Technology & Operational Insights

PRACTICAL NEXT STEP

CCO SERVICE

AI, Automation & Digital Transformation

Automation and intelligent systems, implemented responsibly.

www.clientcareoperations.com/services/ai-automation

FREE PLANNING TOOL · PDF · WORD

AI Readiness Assessment

Seven practical dimensions to examine before introducing AI into a business workflow.

www.clientcareoperations.com/resources#tools

Prefer to talk it through? Discuss an AI use case: www.clientcareoperations.com/contact

Sources and further reading

- [1] **NIST**, *AI Risk Management Framework Playbook*
<https://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook>
- [2] **Office of the Privacy Commissioner of Canada**, *Privacy Guide for Businesses*
https://www.priv.gc.ca/en/privacy-topics/privacy-laws-in-canada/the-personal-information-protection-and-electronic-documents-act-pipeda/pipeda-compliance-help/guide_org/
- [3] **OWASP GenAI Security Project**, *OWASP GenAI LLM Top 10 2026*
<https://genai.owasp.org/resource/owasp-genai-llm-top-10-2026/>
- [4] **Gouvernement du Québec (LégisQuébec)**, *Act respecting the protection of personal information in the private sector (CQLR c P-39.1), sections 3.3 and 12.1*
<https://www.legisquebec.gouv.qc.ca/en/document/cs/P-39.1>

FURTHER READING

- **OWASP GenAI Security Project**, *2025 Top 10 for LLM and Gen AI Applications (previous edition)*
<https://genai.owasp.org/llm-top-10/>

This is an original educational framework inspired by external guidance, not an official NIST assessment or legal/compliance certification. References are ordinary outbound links; no endorsement by, or partnership with, the referenced organizations is implied.