

OSFI's AI warning: waiting is a decision, too

Three investment questions for technology, cyber and risk leaders at Canadian banks and insurers

13 October 2026 | Independent analysis

A security team wants funding to improve how it finds and fixes vulnerabilities. An operations team wants an AI assistant to handle more of a claims workflow. Procurement has a vendor promising both speed and control. All three proposals arrive under the same heading: AI. They need different decisions.

That is a useful way to read the Office of the Superintendent of Financial Institutions' October 8 risk outlook. Alongside warnings about AI-enabled attacks and concentrated technology providers, OSFI identifies reputational risk from delayed adoption. It places particular emphasis on keeping cyber security and vulnerability-identification capabilities current. [1]

For a bank or insurer, simply waiting is therefore a choice with consequences. So is deploying a more capable system before its controls and operating model are ready. This guide offers a way to examine both choices without assuming that every AI proposal deserves funding.

The useful question: Which investment would improve a specific business or security outcome, and what evidence would justify the authority and dependency that come with it?

Separate the three decisions

OSFI's outlook is a supervisory assessment, not a shopping list. It does not establish that a particular model, agent platform or vendor is suitable for your institution. The investment case still has to connect capability to an identifiable need.

Decision	What the discussion should establish
Strengthen defences	Which current weakness would improve, and how would you measure the improvement?
Delegate work	Which actions should the system perform, and which still require a person or another control?
Accept dependency	Which services would become harder to operate, recover or move if the provider changed or failed?

This editorial framework is not an OSFI-prescribed template. A claim about AI capability cannot answer all three questions.

Fund the improvement, then test the claim

Start with the constraint. Is the problem discovering weaknesses, deciding which matter, assigning remediation, or getting fixes into production? Improving one step may leave the real bottleneck untouched.

OSFI's July generative and agentic AI bulletin suggests considering AI to prioritize risks and accelerate patching. It also calls attention to governance, monitoring and controls across the AI lifecycle. Those are sound practices to consider, not a guarantee that a product will improve your process. [2]

Suppose an AI tool produces a better-ranked vulnerability queue. That may be useful. But if the service owner cannot schedule a change, the exposure remains. A credible evaluation would follow the item through triage, ownership, action and verification, rather than ending when the tool produces a convincing recommendation.

Before a pilot, agree on a baseline. Suggested measures include time to a validated remediation decision, missed high-impact issues, false positives and work shifted onto existing staff. Choose measures that fit the proposed use.

For a deferral, name the reason, current exposure and condition that would reopen the decision. A time-limited deferral can be sensible. An indefinite pause with no owner tells the next team very little.

Buy only the autonomy you can operate

Consider an illustrative claims assistant. It can summarize evidence for an adjuster, update a claim record, or initiate a payment. Each step changes what a mistake can do. Success at summarization does not demonstrate readiness to move money.

The July bulletin discusses limits on autonomy, approval checkpoints for high-impact actions, scoped permissions and records of agent activity. It also notes that changes to prompts, models, data sources, tools or third-party services can change system behaviour. [2]

Can the service deliver useful value with narrower permissions? Can the institution review an exception, suspend execution and reconcile unfinished work? Who carries that operational cost?

An initial release might prepare recommendations while leaving payment execution with an established control. A later release could receive more authority after testing. Make that authority a deliberate choice, rather than an accidental consequence of buying the product.

This also makes model governance more useful. E-23 includes AI and machine learning in its model definition, but it does not explicitly name AI agents. A model-risk decision and a decision about what the surrounding system may do answer different questions. Both may matter to the same service. [5]

Price the dependency before it becomes difficult to change

OSFI identifies concentration among frontier-model and cloud providers as a potential source of correlated disruption. Two products can look different to a buyer while depending on the same underlying service. That makes the architecture behind the contract relevant to the investment decision. [1]

B-10 expects institutions to assess third-party risk proportionately, including concentration, subcontracting, continuity and exit. A small, low-risk arrangement does not need the same treatment as a service supporting a critical operation. [3]

For a consequential AI service, ask what it would take to continue without the selected provider. Include the data, integrations, people, alternative capacity and time needed to restore an acceptable service. If the proposed fallback is manual, test the workload people could actually absorb. If it is another vendor, examine the dependencies the two providers share.

Integration, oversight and recovery belong in the cost discussion. A cheaper licence may require a more expensive operating model. Make those assumptions visible before committing.

Use the regulatory dates to sequence work

The revised E-23 model-risk guideline takes effect on **May 1, 2027**. That is a preparation milestone, not the start of all controls over an AI-enabled service. B-13's technology and cyber-risk expectations and B-10's third-party expectations also matter within their applicable scope. [3, 4, 5, 6]

E-21 has a separate implementation sequence. Its letter sets **September 1, 2026** for full adherence, including critical-operation identification, mapping and disruption tolerances, plus a testing methodology and testing begun. Testing across all critical operations is to be completed by **September 1, 2027**. [7]

For a proposed investment, the practical implication is to connect the new capability to work already under way: service ownership, change management, third-party oversight, model review and recovery testing. Creating an isolated AI workstream can leave the same operating question unanswered in several places.

Keep the status of the source clear. Guidelines express supervisory expectations; the July bulletin offers practices institutions can consider. FIFAI II's AGILE framework provides industry perspectives and explicitly is not regulatory guidance. Applicable legal obligations require their own assessment. [2, 8]

Leave the meeting with a decision that can change

Take one proposed AI investment and ask its sponsors to bring three things: the improvement they expect, the authority the service would need, and the dependencies the institution would accept. Then ask what evidence could change their recommendation.

A useful outcome could be to proceed, proceed with conditions, narrow the scope, choose an alternative, or defer. Record the owner, unresolved assumptions and next review point. The aim is to make uncertainty manageable enough to act on, while keeping the decision open to better evidence.

Bring a live decision. If an AI investment or deployment is approaching approval, an independent review can help examine its architecture, cyber exposure, authority and operating assumptions. <https://trustcyber.ca/advisory/>

General information, not legal or regulatory advice. Written for technology, cyber and risk leaders at Canadian federally regulated banks and insurers; applicability elsewhere differs. Examples, proposed measures and investment questions are editorial analysis, not findings about an institution or prescribed OSFI requirements. Sources checked October 13, 2026.

Sources

Primary OSFI publications. Citation numbers refer to this list. Full URLs are visible and clickable.

[1] 2026-2027 Semi-Annual Risk Outlook

<https://www.osfi-bsif.gc.ca/en/about-osfi/reports-publications/annual-risk-outlook-semi-annual-update-fiscal-year-2026-2027>

[2] Generative and Agentic AI technology risk bulletin

<https://www.osfi-bsif.gc.ca/en/risks/technology-cyber-risk-management/technology-risk-bulletin/generative-agentic-artificial-intelligence-implications-technology-cyber-security-operational>

[3] Guideline B-10 - Third-Party Risk Management

<https://www.osfi-bsif.gc.ca/en/guidance/guidance-library/third-party-risk-management-guideline>

[4] Guideline B-13 - Technology and Cyber Risk Management

<https://www.osfi-bsif.gc.ca/en/guidance/guidance-library/technology-cyber-risk-management>

[5] Guideline E-23 - Model Risk Management (2027)

<https://www.osfi-bsif.gc.ca/en/guidance/guidance-library/guideline-e-23-model-risk-management-2027>

[6] Guideline E-23 (2027) release letter and feedback

<https://www.osfi-bsif.gc.ca/en/guidance/guidance-library/guideline-e-23-model-risk-management-2027-letter>

[7] Guideline E-21 release and implementation letter

<https://www.osfi-bsif.gc.ca/en/guidance/guidance-library/operational-risk-management-resilience-letter>

[8] FIFAI II: Adopting an AGILE Framework in Canadian Financial Services

<https://www.osfi-bsif.gc.ca/en/about-osfi/reports-publications/fifai-ii-ai-risks-opportunities-adopting-agile-framework-canadian-financial-services>