HOSMAK
SOLUTIONS
Govern Smarter. Manage Risk Better.

AI & model governance readiness checklist

A self-assessment for vendors and financial services firms preparing for OSFI E-23

When a regulated financial institution assesses your AI or model governance, its reviewer is trying to establish six things. This checklist turns those six questions into a self-assessment you can complete in under an hour. Answer honestly — the point is to find the gaps before your client's questionnaire does.

How to use it: for each item, mark Yes only if you could attach evidence today — a document, a log, a named owner. “We do this informally” is Partly. Then use the scoring guide on the last page.

1 Do you know what your models are?

Why they ask: an institution cannot govern what its vendor has not identified. The inventory is the foundation for every other answer — and the most common gap.

Self-assessment item	Yes	Partly	No
We maintain a written inventory of every model-like component in our product (anything that scores, ranks, predicts, classifies, flags, prices or generates).	☐	☐	☐
Each inventory entry records the model's purpose, input data, and the business decisions its output influences.	☐	☐	☐
Each model is assigned a risk tier based on materiality and impact.	☐	☐	☐
The inventory has a named owner and a defined review cycle.	☐	☐	☐
Third-party and open-source models embedded in our product are in the inventory, not just the ones we built.	☐	☐	☐

Evidence we could attach today: _____

2 Does someone independent check them?

Why they ask: “the team that built it tested it” is not validation. Reviewers look for demonstrable distance between builder and checker, and a record of what the checker found.

Self-assessment item	Yes	Partly	No
Model design, assumptions and limitations are reviewed by someone outside the team that built the model.	☐	☐	☐
That review is documented: scope, findings, and how findings were resolved.	☐	☐	☐
Our highest-tier models receive the deepest review, on a defined cycle.	☐	☐	☐
Material issues found in review are tracked to closure, with dates and owners.	☐	☐	☐

Evidence we could attach today: _____

3 Do you watch them in production?

Why they ask: a model that was sound at deployment can drift. A dashboard nobody reviews is worse than none — it proves you knew what to look at.

Self-assessment item	Yes	Partly	No
Each production model has defined performance metrics and thresholds.	☐	☐	☐
Threshold breaches trigger a defined response, with a named responder.	☐	☐	☐
Monitoring results are reviewed on a set cadence, and the reviews are logged.	☐	☐	☐
We can show at least one example of a monitoring alert that led to action.	☐	☐	☐

Evidence we could attach today: _____

4 Do you control changes?

Why they ask: an unannounced model change lands inside the client's control environment. Institutions need to know how changes reach production and how they will hear about the ones that affect them.

Self-assessment item	Yes	Partly	No
Movement of a model from development to production follows a documented, approved path.	☐	☐	☐
Retraining and material model changes require documented approval before release.	☐	☐	☐
We notify affected clients of changes that could alter model behaviour they rely on — beyond a line in the release notes.	☐	☐	☐
We can reconstruct which model version was in production on a given date.	☐	☐	☐

Evidence we could attach today: _____

5 Is anyone accountable?

Why they ask: committees diffuse accountability; risk functions look for a person. “The ML team” is not an answer a reviewer can write down.

Self-assessment item	Yes	Partly	No
Every model in the inventory has a named individual owner.	☐	☐	☐
Ownership is written into a role description, not held informally.	☐	☐	☐
There is a defined escalation path when a model misbehaves, and staff know it.	☐	☐	☐
Someone at leadership level is accountable for model governance overall.	☐	☐	☐

Evidence we could attach today: _____

6 What are the limits?

Why they ask: to a risk reviewer, “no known limitations” reads as “no one has looked.” Articulating limits scores better than claiming none.

Self-assessment item	Yes	Partly	No
Each model's documented profile includes known limitations and failure modes.	☐	☐	☐
We define the conditions under which the model's output should not be relied on without human review.	☐	☐	☐
Data constraints (coverage, bias, recency) are documented per model.	☐	☐	☐
Client-facing documentation states what the model should and should not be used for.	☐	☐	☐

Evidence we could attach today: _____

Scoring your self-assessment

Count your answers across all 26 items.

Mostly Yes	You are ahead of most vendors. The remaining work is assembling your evidence into a two-to-four page governance summary that can be attached to any client questionnaire — the artefact that shortens due diligence cycles most.
Mixed Yes / Partly	Your practices are probably sounder than your paperwork. The gap is documentation and evidence, which is the fastest kind to close — typically six to ten weeks of structured effort. Start with the inventory (section 1); everything else hangs off it.
Mostly Partly / No	You have a governance build ahead of you, and the calm version of that build takes a quarter. Done under the pressure of a live questionnaire with a fourteen-day turnaround, it reads as what it is. The clock that matters is your clients' — regulated institutions need their third-party model evidence in place before OSFI Guideline E-23 takes effect on 1 May 2027, so their questions arrive well before that date.

Three things to do this quarter

1. **Build the inventory.** A week or two of honest work, and the foundation for every other answer in this checklist.

2. **Write the governance summary.** Two to four pages answering the six questions above, ready to attach to any questionnaire before one arrives.

3. **Pressure-test both.** Read them as the institution's reviewer would. If a claim has no evidence behind it, it is a gap, not an answer.

Want a second pair of eyes from someone who has sat on the reviewing side?
Hosmak Solutions advises financial services firms and their vendors on AI governance, model risk management and technology GRC — led by 25 years in financial services technology risk, including Big 4 advisory and nearly a decade building enterprise technology risk governance at one of Canada's largest banks.

Book a 30-minute consultation: hosmaksolutions.ca · hello@hosmaksolutions.ca