Supervising AI in Finance

From Policy Principles to Practical Controls

Strengthening trust, control and accountability in financial AI

Presented by

Dr Dharma Basgeet

General Manager

Harel Mallac Technologies

RCE-OECD Virtual Training | 26 May 2026

E: dharma.basgeet@harelmallac.com

M: +23059421118

harel mallac
technologies

AI Is No Longer a Future Issue in Finance

AI is already embedded in financial operations — making trust, control and accountability central to supervision.

AI is already inside your institution

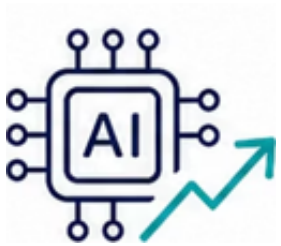

***75%** of financial firms are already using or testing AI*

AI is no longer in the lab — it is already influencing credit decisions, fraud detection, compliance monitoring and customer interactions.

The risks are real — and governance is still catching up

Only 12% of firms have formal AI risk governance frameworks in place

AI risk is widening — from explainability and bias to data quality, cyber exposure and third-party dependency.

Supervisors are raising the bar — fast

Global expectations are moving toward stronger AI governance, resilience and accountability.

The shift is clear: from AI principles to evidenced governance, testing, monitoring and accountability.

Supervisory Shift → AI experimentation → AI adoption → AI operational dependency → AI supervision

Reference context: OECD | FSB | BIS | EU AI Act | ACA Group / NSCP 2024 AI Benchmarking Survey

AI Is Moving Across the Financial Value Chain

AI is becoming embedded across customer engagement, decisioning, risk controls and operations

Customer Engagement

Examples

Chatbots, onboarding, personalised support

Supervisory Lens

Customer protection & escalation

Financial Decisioning

Examples

Credit scoring, underwriting, pricing

Supervisory Lens

Fairness, bias & explainability

Risk & Financial Crime

Examples

Fraud alerts, AML, anomaly detection

Supervisory Lens

Model performance & model drift

Compliance & Governance

Examples

Regulatory reporting, surveillance, policy review

Supervisory Lens

Auditability, evidence & accountability

Operations & Resilience

Examples

Automation, document processing, productivity tools

Supervisory Lens

Resilience, data security & third-party dependency

Supervisory Implication

AI use cases must be mapped, risk-rated and controlled across the full financial value chain.

- Different use cases.
- Different risks.
- Different supervisory focus.

The Risk Is Not the AI — It Is What the AI Influences

Supervision should be proportionate to the level of decision influence, customer impact, control dependency and accountability risk.

AI Influences	Supervisory Questions
Customer outcomes	Could the AI mislead, exclude or unfairly treat customers?
Financial decisions	Can the decision be explained, challenged and overridden?
Risk controls	Is the model accurate, monitored and resilient to drift?
Regulatory evidence	Can the institution prove what happened, why and who approved it?
Operational dependency	Who is accountable if the AI, data pipeline or third-party provider fails?

Supervision should follow impact – NOT hype

Supervisory Implication

The higher the customer impact, decision influence or control dependency, the stronger the governance and evidence required

Seven Practical Challenges for AI Supervision

Effective supervision requires looking beyond the algorithm — into data, controls, monitoring, third parties and accountability.

OUTCOME RISK

1 Explainability

Can the outcome be explained?

2 Bias & Fairness

Could customers be unfairly treated?

CONTROL RISK

3 Data Quality & Lineage

Is the data trusted and traceable?

4 Model Risk & Drift

Is performance monitored over time?

OPERATIONAL RISK

5 Cybersecurity & Misuse

Can the AI or data pipeline be manipulated?

6 Third-Party Dependency

Who controls the model and platform

7 Accountability & Governance

Who owns the decision and the failure?

- *Clear ownership of AI outcomes*
- *Approval of changes and models*
- *Response and remediation when something goes wrong*

Supervisory implication: AI supervision must test evidence across the full lifecycle – from design and data to deployment, monitoring, change and accountability.

Case Study: GenAI Customer Assistant in Finance

A simple-looking customer tool can become a supervisory risk point when it touches data, decisions, vendors and customer outcomes.

Scenario

A financial institution deploys a GenAI-powered assistant to support customer onboarding, product questions and loan eligibility guidance.

The assistant connects to customer data, product rules, internal workflow systems and a third-party AI/cloud platform.

How The Assistant Works

1	2	3	4	5	6
Customer Query	**Gen AI assistant**	**Customer data & product rules**	**Internal systems**	**Third party AI / Cloud platform**	**Response to customer / Human escalation**
Customer asks a question via chat / voice / web / mobile	AI interpretes the query and generates the response	Assistant accesses customer profile, KYC, products, policies & rules	Connects to CRM, loan systems, workflows and knowledge base	Model inference hosting or other AI services provided by a third party	Response is sent back to customer or escalated to a human agent

Pre-Go-Live Challenge: What Must Be Proven?

Customer protection
Could the response mislead or unfairly treat the customer?

Data governance
What data is accessed, stored or shared?

Auditability
Are prompts, responses and decisions logged?

Human oversight
When does the AI escalate to a human?

Accountability
Who is accountable if something goes wrong?

Supervisory implication: Customer-facing GenAI requires controls over data, outputs, escalation, auditability, vendors and accountability before deployment.

From AI Use Case to Go-Live Readiness

High-impact AI should not go live until key controls are evidenced.

AI Use Case — **Go-Live Readiness Gates** — **Go / No-Go**

1 Purpose & Risk Classification

What is the AI used for and how material is the impact?

Evidence: Inventory, risk rating, business owner

2 Governance & Approval

Who approved the use case and under what risk appetite?

Evidence: Approval record, governance forum, policy alignment

3 Data Controls

What data is used, where is it processed and is it authorised?

Evidence: Lineage, legal basis, access and retention controls

4 Model Testing & Validation

Has AI been tested for performance, bias, drift, and explainability?

Evidence: Validation, test results, bias assessment, thresholds

5 Cyber & Third-Party Risk

Could the AI environment be compromised?

Evidence: Security review, vendor due diligence, contractual safeguards, incident plan

6 Human Oversight & Escalation

When does a human review, challenge or override the output?

Evidence: Escalation rules, override process, exception handling

7 Auditability & Accountability

Is there a clear audit trail and owner?

Evidence: Logs, audit trail, ownership matrix, remediation

Go-live principle: No high-impact AI use case should go live without clear ownership, tested controls, audit evidence and escalation mechanisms.

From Approval to Continuous AI Control

AI governance must operate as a lifecycle control system — with clear ownership, evidence, monitoring and remediation.

5 *Does the institution detect issues and fix them over time?*

Evidence: Monitoring dashboards, audit logs, incidents, remediation plan

4 *When do humans review, challenge or override AI outputs?*

Evidence: Human-in-the-loop rules, escalation paths, exception handling

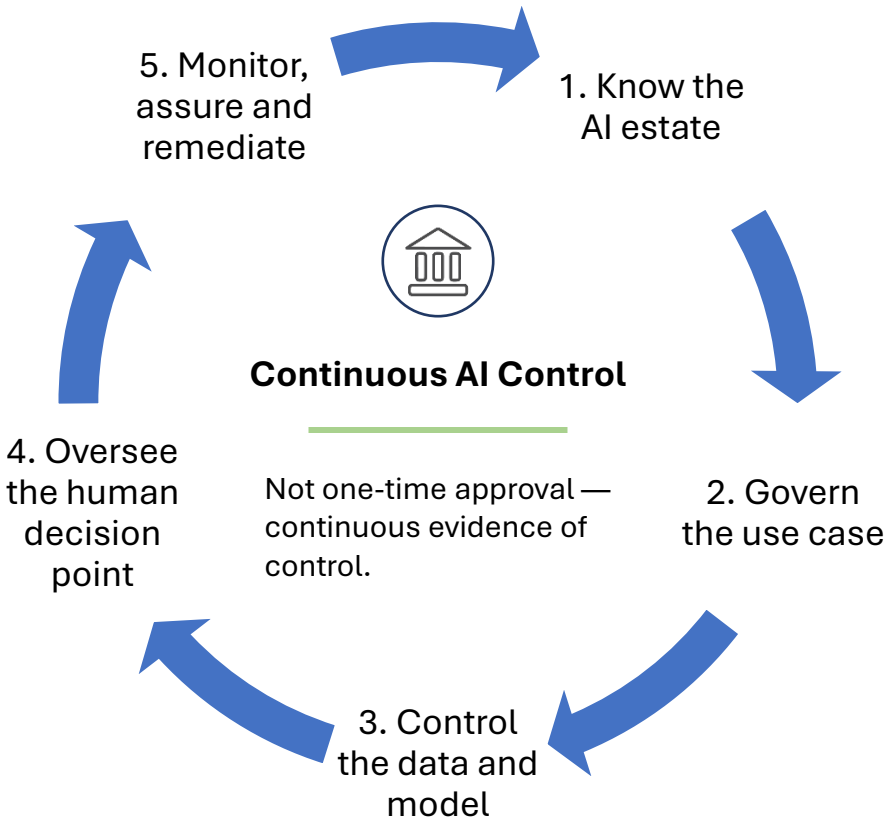

1 *Where is AI being used and how material is the risk?*

Evidence: AI inventory, risk tiering, business owner, purpose

2 *Who approved it and under what risk appetite?*

Evidence: Governance forum, approval record, RACI, policy alignment

3 *Are data, model performance, fairness and explainability tested?*

Evidence: Data lineage, validation, bias testing, drift thresholds

Supervisory implication: The control question is not only "was the AI approved?" — it is "does the institution remain in control after deployment?"

AI Accountability Must Be Enterprise-Wide

Supervision should test whether accountability is clear across the Board, management, business, risk, technology and third-party providers.

	WHAT THEY OWN	SUPERVISORY QUESTION
1	**Board / Board Risk Committee** Sets risk appetite, oversees high-impact AI use cases and holds management to account.	**Is the Board informed of material AI risks and comfortable with the risk appetite?**
2	**Executive Management** Owns AI strategy, operating model and cross-functional governance.	**Who is accountable for AI outcomes across the institution?**
3	**Business Owner** Owns the AI use case, customer impact and business decision logic.	**Does the AI support the intended business outcome safely and fairly?**
4	**Risk, Compliance & Legal** Challenges risk classification, fairness, regulatory impact and evidence.	**Are controls sufficient and regulatory obligations met?**
5	**Technology, Data & Cybersecurity** Controls architecture, data, model environment, access, security and monitoring.	**Is the AI environment secure, reliable and monitored?**
6	**Third Parties / Vendors** Provide model, cloud or platform services under controlled arrangements.	**Are vendor risks contractually and operationally controlled?**

Key shift:

AI governance is not an IT control; it is an enterprise accountability model.

Supervisory implication: Accountability cannot be outsourced to the model, the vendor or the technology team — ownership must be explicit, senior and evidenced.

Mauritius: Practical, Proportionate AI Supervision

Enabling responsible AI adoption through risk-based oversight, control evidence and ecosystem readiness.

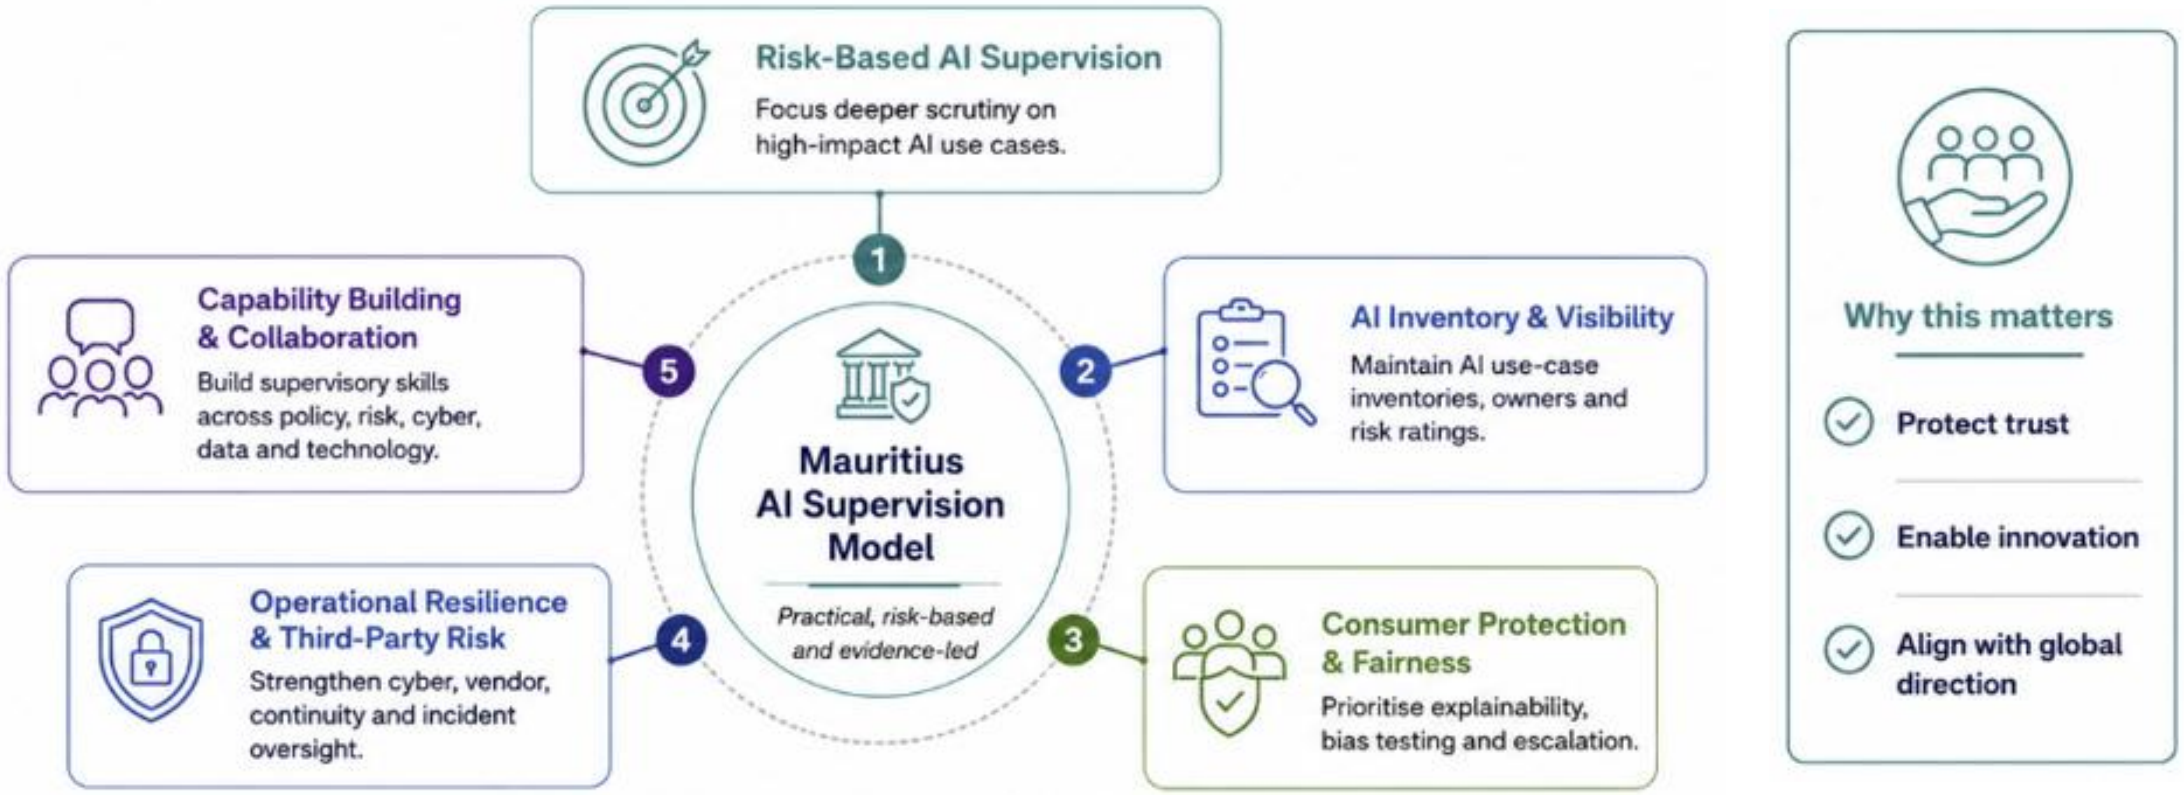

Supervisory Implication: Mauritius has an opportunity to protect trust without slowing innovation — through practical, proportionate and evidence-led AI supervision.

Trust Is the Licence to Scale AI

In finance, AI will only scale when institutions can prove it is governed, controlled, monitored and accountable.

1

Supervision follows impact

The deeper AI influences customers, decisions and controls, the stronger the scrutiny should be.

2

Evidence replaces intention

Institutions must demonstrate governance, testing, monitoring, auditability and accountability — not only declare responsible AI principles.

3

Trust enables scale

Practical, proportionate supervision protects confidence while allowing responsible AI adoption to grow.

The objective is not to slow AI down — it is to make AI safe enough, trusted enough and accountable enough to scale

Let’s Continue the Conversation

We would be pleased to engage further on practical, trusted and responsible AI adoption in financial services.

Head Office

Harel Mallac Technologies Mauritius

1st Floor, Block A, Phoenix Central

Post Pitt, Phoenix 73524

T: (+230) 207 3300

Burundi

10, Avenue du Commerce, BP 723

Bujumbura, Burundi

T: (+257) 22 27 71 93

Madagascar

Zones Industrielles, Groupe FILATEX Batiment D9

Ankadimbahoaka Route Nationale N°7

Antananarivo, Madagascar

T: (+261) 20 22 543 21

Rwanda

4th Floor Muhima Investment Company (MIC) Building

KN7 Avenue, Nyarugenge

Kigali, Rwanda

T: (+250) 252 57 05 53

---

hmtinfo@harelmallac.com