Technical training on

Supervision of AI in finance

Virtual
Date: Tuesday, 26 May 2026
Time: 10.00-12.00hrs (Mauritius Time)

Session 1: From Policy Principles to Practical Controls

Question 1:

As AI adoption in finance accelerates, how can supervisors move beyond high-level AI principles and assess whether a financial institution's AI governance is actually effective in practice — particularly in areas such as model validation, explainability, data quality, human oversight, third-party reliance, and accountability for AI-driven decisions?

Answer 1:

I would start with a proportionate and risk-based approach. As I mentioned in the presentation, we should not treat every AI use case in the same way. An internal tool used to summarise meeting notes does not carry the same risk as an AI tool supporting customer onboarding, credit eligibility or fraud monitoring.

So the first step is visibility. Institutions should maintain an AI inventory, classify use cases by impact, and assign clear business ownership. Once we know where AI is being used and what it influences, we can decide where deeper scrutiny is required.

For high-impact AI — where it affects customers, credit decisions, anti-money laundering, fraud monitoring, regulatory reporting or operational resilience — supervisors should expect stronger evidence: board visibility, governance approval, model validation, explainability, data protection controls, human oversight, audit logs and third-party due diligence.

The key is practicality. Lower-risk tools should be managed proportionately, but high-impact AI should be subject to clear control expectations before and after deployment. In simple terms: supervision should follow impact, accountability should be explicit, and responsible AI must be evidenced — not only declared.

Question 2:

n many jurisdictions, including smaller financial centres, AI is entering finance through vendor platforms, cloud-based tools, compliance systems, credit scoring, fraud monitoring and customer service. How should supervisors build an effective AI supervisory framework that is proportionate, risk-based and practical, while still ensuring board accountability, model validation, explainability, data protection, consumer fairness and control over third-party dependencies?

Answer 2:

I would start with a proportionate and risk-based approach. As I mentioned in the presentation, we should not treat every AI use case in the same way. An internal tool used to summarise meeting notes does not carry the same risk as an AI tool supporting customer onboarding, credit eligibility or fraud monitoring.

So the first step is visibility. Institutions should maintain an AI inventory, classify use cases by impact, and assign clear business ownership. Once we know where AI is being used and what it influences, we can decide where deeper scrutiny is required.
For high-impact AI — where it affects customers, credit decisions, AML, fraud monitoring, regulatory reporting or operational resilience — supervisors should expect stronger evidence: board visibility, governance approval, model validation, explainability, data protection controls, human oversight, audit logs and third-party due diligence.

The key is practicality.

Question 3:

Hi, my question refers to AI influencing Customer Engagement.

How far or to what extent can deepfake AI affect customer identification or verification process and according to you what can be done to prevent customer impersonation and even identity fraud.

Answer 3:

This is a very relevant question because deepfakes attack one of the most important control points in finance: trust in customer identity.

In digital onboarding or remote verification, deepfakes can imitate a customer's face, voice, document or live interaction. The risk is that a fraudster may appear to pass verification even though the real customer is not present.

So institutions should not rely on one control, such as facial recognition alone. They need layered identity assurance: liveness detection, document verification against trusted sources, device and behaviour risk signals, human escalation for higher-risk cases, and audit logs to reconstruct what happened.

The key point is this: deepfake risk does not mean digital onboarding should stop. It means identity verification must become stronger — combining technology, trusted data, human review and auditability.

This links directly to my presentation: when AI touches customer engagement and onboarding, the question is not only whether the tool works; it is whether identity, data, controls and accountability are protected.

Question 4:

with generative AI system in FI, how to ensure auditability as Gen AI uses neuralpathways versus logical pathways

Answer 4:

With GenAI, auditability cannot depend only on explaining the internal neural pathway. The approach is to audit the control environment: what was asked, what data was used, what response was produced, who approved it, when a human intervened and how the output was monitored (all 7 key points I covered during my presentation). In finance, the key is not only to understand the model internally, but to ensure the institution can reconstruct what happened, evidence the controls and assign accountability.

Question 5:

What organisational changes are necessary for regulators to become AI-ready institutions?

Answer 5:

Touched on this in the presentation and believe an AI-ready regulator is not one that tries to inspect every algorithm line by line. It is one that knows where AI is being used, focuses on the use cases that matter most, challenges evidence, and supervises AI in a risk-based, proportionate and practical way.

Question 6:

Should regulators establish a dedicated AI Governance Committee at board level, or integrate AI oversight into existing Risk or Technology Committees?

Answer 6:

That is a very relevant follow-up, and it links directly to the point I made on enterprise-wide AI accountability. In my presentation, I highlighted that AI accountability cannot sit only with technology teams or vendors. It must be explicit, senior and evidenced.

There is no one-size-fits-all answer. A dedicated AI governance forum may be useful at the early stage to build focus, consistency and capability.

But over time, AI oversight should be embedded into existing governance structures, especially risk, technology, data, cybersecurity, conduct and operational resilience forums. Otherwise, AI risks becoming a separate side topic rather than part of mainstream supervision.

The key is not the committee name. The key is clear ownership, escalation, evidence review, and visibility at senior management and board level.

So, to connect it back to my slide on accountability: AI governance must be enterprise-wide. A dedicated forum can help create momentum, but lasting oversight must be integrated into the institution's broader risk and governance model.

Question 7:

is our regulatory environment mature for gen AI in the client onboarding cycle. We see abroad many FIs use gen AI in onboarding. Are we not lagging behind while still sorting the GRC of AI locally while internationally FIs are being more competitive with use of AI

Answer 7:

I would not frame it as Mauritius lagging. I would frame it as a global implementation challenge.

Many jurisdictions are asking the same question: how do we use generative artificial intelligence in onboarding while protecting identity, data, customer due diligence and financial crime controls?

The answer is controlled adoption. Generative artificial intelligence can help with document summarisation, customer guidance, checklist validation and operational efficiency. But final decisions, high-risk cases, identity verification and exceptions must remain subject to strong controls and human oversight.

So we should move fast, but safely. Do not block generative artificial intelligence in onboarding, but do not treat it as a simple chatbot either. Treat it as a high-impact customer journey use case that requires evidence, governance and accountability.

Question 8:

Is there regulatory framework in Mauritius that determine the use of AI tools and supervision, especially for monitoring financial institution?

Answer 8:

I would separate two things: the legal design of a national framework and the practical supervisory approach.

The legal framework is a policy decision for the relevant authorities. From a practical supervisory perspective, the key is that AI oversight should be risk-based, proportionate and evidence-led.

Whether through guidance, standards, supervisory expectations or future legal frameworks, the important point is that high-impact AI use cases should have clear ownership, testing, data controls, human oversight, auditability and accountability.

So my focus would be less on the label of the framework and more on whether institutions can demonstrate that AI is governed, controlled, monitored and accountable in practice.

Question 9:

Currently we have different forms of regulations/framework at level of different institutions .. would not an AI Act be more useful in the future ?

Answer 9:

That is an important policy question, and I would be careful not to comment on the legal design of a national framework.

From a practical supervisory perspective, the real question is what we want the framework to achieve. Whether it is through an Act, guidance, standards or supervisory expectations, the

outcome should be the same: AI use should be risk-based, proportionate, transparent, accountable and evidence-led.

An AI Act may be useful in the future if the policy objective is to create a common national baseline. But it should avoid becoming a one-size-fits-all constraint that slows responsible innovation.

For me, the priority is that high-impact AI use cases — those affecting customers, financial decisions, regulatory evidence, financial crime controls or operational resilience — must be subject to stronger governance, testing, human oversight, third-party controls and auditability.

So the key point is this: the form of the framework is a policy decision; the practical control objective is clear — responsible AI must be demonstrated, not only declared.

Question 10:

How does a risk based supervisory approach scale obligations to a small fintech vs a large insurer?

Answer 10:

A risk-based approach does not mean weaker standards for smaller firms. It means supervision is scaled to the size, complexity and risk of the activity.

A small fintech using AI for a low-risk internal process may need lighter evidence: ownership, basic controls, testing and escalation.

A large insurer using AI for underwriting, pricing or claims decisions should face stronger expectations: board visibility, model validation, explainability, bias testing, audit trails and ongoing monitoring.

So the principle is simple: obligations should follow the impact of the AI use case — not just the size of the institution.

Question 11:

the fsc currently has a guidance notes which is non-binding. Since mauritius is still baby journeying compared to bigger jurisdiction, is it not the right time to have an AI act like EU to ensure responsible use of AI in a phased approach?

Answer 11:

I would avoid framing Mauritius as behind. Many jurisdictions, including larger ones, are still working through how to translate artificial intelligence principles into practical supervision.

The existence of guidance notes is a useful starting point because it helps establish direction. The question of whether Mauritius should move toward an Artificial Intelligence Act is a policy decision for the relevant authorities.

From a practical supervisory perspective, what matters is that the approach remains phased, proportionate and risk-based.

A phased approach could make sense: start with high-impact AI use cases, strengthen evidence expectations, build institutional capability, and then decide whether further legal or regulatory instruments are required.

So my answer would be: guidance is a positive starting point, and the next step is disciplined implementation. Whether that evolves into an Act should depend on policy objectives, market maturity and the level of risk that needs to be addressed.

Question 12:

How often should the data be reviewed to ensure that the AI outcome is not bias/ performing more efficiently?

Answer 12:

There is no single fixed frequency that works for every AI model. The review should be risk-based.

For high-impact AI — for example credit, pricing, onboarding, fraud, anti-money laundering or claims decisions — data and model performance should be monitored continuously where possible, with formal reviews at defined intervals or whenever there is a material change.

For lower-risk AI, periodic review may be sufficient.

The key is not only frequency. Institutions should also review the model when customer behaviour changes, market conditions change, data sources change, the model is updated, or unusual outcomes appear.

In simple terms: monitor continuously where the risk is high, review formally at agreed intervals, and trigger an immediate review when the data, model or outcomes change.

Question 13:

What are the risks to the local economy and local economic activity when customers use AI to make their financial decisions and most especially when they are subject to AI generated marketing of goods online? Is it not a threat to domestic activity and to local businesses? How can we address this issue?

Answer 13:

I would not position AI only as a threat. Like any powerful tool, AI creates risk if it is unmanaged, but it also creates opportunity if it is governed and used responsibly.

The underlying issue is not entirely new. For many years, businesses have used data, profiling, targeted advertising and recommendation engines to influence consumer behaviour. What AI changes is the speed, scale and personalisation of that influence.

So yes, there is a risk that AI-generated marketing and digital platforms can influence customer choices, create unfair outcomes, or reduce visibility for smaller local businesses. But the answer should not be to block AI. The answer is to manage the risk and help local players become AI-ready.

For Mauritius, the right response should be practical: ensure transparency when AI is used, protect consumers from misleading recommendations, promote responsible use of data, encourage fair digital competition, and build capability among local businesses and financial institutions.

Session 2: AI agents: the new finance person

Question 1:

Agentic AI introduces a new supervisory challenge because the system may not only analyse information but also initiate actions across workflows. How should regulators define the point at which AI autonomy becomes material enough to require enhanced governance, independent validation, real-time monitoring, human override, and possibly prior regulatory notification?

Answer 1:

In the case some critical information is not analysed by an agent, this would usually denote a very badly evaluated AI development process with very low level safeguards. Hence, the probability of this happening should be considered as low. For regulators to determine which AI system requires human oversight, the risk tiering approach is usually applied, where the impact of the AI system determines linearly the level of human oversight.

Question 2:

Realistically, which finance jobs in Mauritius do AI agents replace in the next five years, and which ones become valuable??

Answer 2:

I cannot give any precise answer to this as too many variables involved. However, Agentic AI are connected with jobs which are very much data driven and rule based. However, the involvement of humans would still be needed to oversee, validate, curate inputs etc.

Question 3:

When an AI model's data lineage cannot be reconstructed after a regulatory challenge whether due to a third-party feed gap, unversioned feature engineering, or missing inference logs; Should the institution treat this as an Operational Risk event, a Model Risk event, or a Data Governance failure

Answer 3:

It won't be a data governance failure as the issues have been detected. However, these are risks and most probably could be categorized as model risks, which is typically a subcategory of operational risk.

Question 4:

Deployment of advanced AI tools carry elements of risk - do you think we should have an AI Licensing / certification system?

Answer 4:

Most probably more of auto-regulation due to lack of capacity from government regulators.

Question 5:

Tech/IT will always be ahead of legal / regulatory framework as there will be no precedence. Grateful for input / guidance on same, thank you

Answer 5:

Yes, and there could be different ways to tackle this. Firstly, the laws could be targeted towards the consequences of technology, and thus the state of the technology does not really matter. E.g, data privacy is upheld as a human right, and thus any technology affecting data privacy would need to be compliant by DPA. Secondly, and I see a major global push towards this currently, is to ensure that lawmakers/ judicial systems possess a much higher capacity to appreciate developments/trends in technology. Hence, this could reduce the mentioned gap by coming up with regulations quite quickly after launch of new technology.

Question 6:

Does Mauritius need sovereign compute infrastructure (national cloud / GPU clusters), or is hybrid cloud sufficient?

Answer 6:

Answered verbally. To summarize, there are use cases requiring both. Sovereign compute would be more applicable to strategic national use cases, thus driven by public sector AI adoption largely. For private sector, having private compute might depend upon updates in data protection laws and evolution of cost of GPUs, but right now it is needed in some use cases, but most use cases might still operate in the hybrid model.

Question 7:

What about assessing the business case before implementing agentic AI? Is it a given?

Answer 7:

The answer is definitely that there should be a clear business case. AI is just a tool helping an organization, so should be aligned to strategic vision.

Question 8:

Hello. How far is Mauritius today ready to incorporate Agentic AI in the financial sector?

Answer 8:

From my point of view, most of the big players could transition quite quickly towards agentic AI; they have some proper data capability, most already have some narrow AI experience and the operational lure can be high. The major issues would be typically processed based, business re-engineering, compliance with laws etc.