



Sound Practices for Responsible Adoption of Artificial Intelligence (AI): Consultation report

Response to Consultation

Nedbank

1. Do you agree with the benefits and risks of AI adoption by financial institutions described in this report? Are there any substantive benefits or risks not covered?

Yes – We agree with the benefits and risks of AI adoption described in the report. However, while the report recognises the importance of cloud infrastructure and the need to consider commercial, computing and energy costs when selecting AI solutions, additional discussion on the scalability and sustainability implications of token consumption, cloud usage and computational costs associated with AI adoption may further strengthen the report. As financial institutions increasingly adopt GenAI and foundation models, these factors may influence the long-term feasibility, operating model and cost-effectiveness of AI deployments.

2. Are the sound practices sufficiently comprehensive and clear to enable financial institutions' responsible AI adoption?

Yes - The proposed sound practices are comprehensive, practical, and appropriately risk-based, providing a strong foundation for responsible AI adoption across financial institutions. The lifecycle-oriented framework, proportionality principle, governance expectations, cyber risk controls, and coverage of GenAI and agentic AI are particularly valuable.

However, the framework would benefit from additional guidance in several areas:

- (i) non-human identity governance for AI agents and autonomous systems;
- (ii) accountability and non-repudiation through immutable audit mechanisms;
- (iii) BCBS239-aligned data, model, prompt and decision lineage;
- (iv) AI supply chain security and model provenance;
- (v) stronger privacy-by-design requirements;
- (vi) AI-specific operational resilience expectations; and

(vii) more explicit controls for high-impact financial-sector use cases such as credit, AML, fraud, trading and regulatory reporting, and

(viii) stronger articulation of AI ethics and responsible use principles, including fairness, bias management, explainability, human oversight, transparency, and protection against unintended customer harm.

3. Do the sound practices strike an appropriate balance between managing risks relating to all forms of AI, and addressing some of the risks relating to emerging and new complex forms of AI, such as GenAI and agentic AI?

Yes – The sound practices generally strike an appropriate balance between managing risks across all forms of AI while recognising the additional risks associated with more advanced technologies such as GenAI and agentic AI. The report appropriately adopts a risk-based and technology-neutral approach, while also including specific considerations for GenAI and agentic AI where relevant. However, the report could be further strengthened through additional practical guidance on AI-specific risk management practices, including the integration of AI within existing model risk management frameworks, accountability for bias and fairness validation, data sovereignty considerations, privacy risk management, the use of quantitative AI risk metrics and KRIs, and managing dependencies across extended AI supply chains, including fourth-party providers.

4. Are the sound practices sufficiently flexible to accommodate and address newer types of AI and responsible adoption over time?

Yes – The sound practices are sufficiently flexible to accommodate evolving AI technologies. The report expressly recognises that AI continues to develop rapidly, that the sound practices are not exhaustive, and that they may need to evolve over time. The principles-based, technology-neutral and risk-based approach, together with the emphasis on proportionality and organisational adaptability, should enable financial institutions to apply the sound practices to emerging AI technologies as they mature.

5. Do the case studies in this report sufficiently highlight how different types of financial institutions can benefit from responsible AI adoption? Are there additional case studies for inclusion in the report? If so, please provide such case studies, particularly for nonbanks.

Yes – However, the following case studies could be added:

(i) Case Study: Asset Management – Investment Research and Portfolio Monitoring

(ii) Case Study: Payment Service Provider – Real-Time Fraud Detection - A payment service provider deployed machine learning and GenAI technologies to identify real-time payment fraud across card, wallet, and account-to-account payment channels.

(iii) Case Study: Financial Market Infrastructure (FMI) - A central securities depository deployed AI-enabled anomaly detection capabilities to identify settlement failures, reconciliation breaks, and emerging operational risks.

6. Do the case studies in this report provide actionable insights for financial institutions in their responsible AI adoption?

Yes - The case studies generally provide actionable insights and are one of the strongest elements of the report. They move beyond theoretical principles and demonstrate how financial institutions have operationalised governance, human oversight, testing, documentation, risk management, and deployment controls in real-world scenarios.

7. Are the definitions in the glossary clear and aligned with industry sound practices, including recent developments in AI?

No - The document extensively discusses AI governance but does not provide a dedicated glossary definition. Further to this, the report discusses AI agents, dynamic Identity and Access Management (IAM), agent identifiers, and synthetic employees elsewhere in the document. However, the glossary does not define: Non-human identities, Machine identities. Agent identities and Workload identities.

For example, AI Agent Identity / Non-Human Identity: A unique digital identity assigned to an AI agent, model, system, service, or workload that enables authentication, authorization, accountability, auditability, and lifecycle management.

8. Are there any comments you would like to make on this report? Please fill in.

The report provides a useful principles-based framework for the responsible adoption of AI, including guidance on governance, risk management, AI lifecycle considerations, third-party risk and emerging technologies such as GenAI and agentic AI. Additional guidance may be beneficial on AI lifecycle management, cloud-specific control considerations within shared responsibility models, and operational resilience testing for AI-enabled processes and third-party dependencies.