DREW & NAPIER

MAS Guidelines on Artificial Intelligence Risk Management: Key Requirements and Implications for Financial Institutions

14 October 2026

LEGAL UPDATE

In this Guide

On 7 October 2026, the Monetary Authority of Singapore ("**MAS**") issued its Guidelines on Artificial Intelligence Risk Management (the "**Guidelines**"), setting out MAS's supervisory expectations for the responsible adoption and use of artificial intelligence ("**AI**") by financial institutions ("**FIs**").

The Guidelines address oversight, key AI risk management systems, policies and procedures, AI life cycle controls, and AI capability and capacity. The Guidelines apply to all FIs in a manner proportionate to their size and risk profile.

The Guidelines take effect on 7 October 2027.

In this Update, we consider:

(a) the scope of the Guidelines and their treatment of AI risks and proportionality;
(b) the principal supervisory expectations for AI oversight, inventories, materiality assessment and life cycle controls;
(c) the relationship between the Guidelines and existing financial-sector and AI governance frameworks; and
(d) the implications for FIs.

03 INTRODUCTION

03 THE SCOPE OF THE GUIDELINES

05 KEY SUPERVISORY EXPECTATIONS

10 RELATIONSHIP WITH EXISTING FRAMEWORKS

11 IMPLICATIONS FOR FIs

11 CONCLUSION

Introduction

On 7 October 2026, the Monetary Authority of Singapore ("MAS") issued its Guidelines on Artificial Intelligence Risk Management (the "Guidelines"), setting out MAS's supervisory expectations for the responsible adoption and use of artificial intelligence ("AI") by financial institutions ("FIs").

The Guidelines address oversight, key AI risk management systems, policies and procedures, AI life cycle controls, and AI capability and capacity. They complement existing MAS legislation and guidance, including the principles on Fairness, Ethics, Accountability and Transparency ("FEAT"), and should be read with other applicable guidance and industry standards.

The Guidelines apply to all FIs in a manner proportionate to their size and risk profile. They apply on a group basis for locally incorporated FIs that are subject to consolidated supervision by MAS or own critical information infrastructure. A FI that is part of a global group may leverage group AI risk management frameworks where those frameworks meet the expectations in the Guidelines.

The Guidelines cover Generative AI and newer developments such as AI agents. They take effect on 7 October 2027. FIs may meet the expectations in Sections 3 and 4 of the Guidelines from 7 October 2027, and the expectations in Sections 5 and 6 of the Guidelines by 7 October 2028.

In this Update, we consider:

(1) the scope of the Guidelines and their treatment of AI risks and proportionality;

(2) the principal supervisory expectations for AI oversight, inventories, materiality assessment and life cycle controls;

(3) the relationship between the Guidelines and existing financial-sector and AI governance frameworks; and

(4) the implications for FIs.

1. The scope of the Guidelines

1.1 Definition of AI

For the purposes of the Guidelines, AI refers to machine-based systems or models that derive outputs through learned premises, such as the data or inputs they receive. The outputs may include estimates, predictions, content, summaries, recommendations or decisions that may influence physical or virtual environments.

The Guidelines distinguish between a model, a system and a use case. A model converts assumptions and input data into outputs; a system comprises one or more models and other machine-based components; and a use case is the specific real-world context in which a model or system is applied. Tools whose outputs are solely based on predefined rules, explicit manually specified formulae or fixed programmed instructions, including ordinary robotic process automation, are not ordinarily AI for these purposes.

1.2 Risks identified

MAS observes that AI-related risks are often an extension of existing risk categories. These include financial risk, operational risk, conduct risk, financial crime risk and reputational risk. The probabilistic and complex nature of AI can create greater uncertainty, unexpected behaviour and difficulty in identifying or explaining inaccurate or biased outputs.

Generative AI may amplify these risks because its unstructured inputs and outputs, opaque training data and limited established testing techniques can make its behaviour harder to evaluate and explain. The Guidelines also identify related security, privacy, legal and intellectual property, third-party, operational and human-factor risks, including shadow AI, automation bias and skill degradation.

AI agents may amplify these concerns further because they may have greater autonomy and access to tools. An agent could execute unauthorised or erroneous actions, or a compromised agent could exfiltrate sensitive data or execute malicious commands at scale.

1.3 Proportionality and basic AI governance policies

The Guidelines require FIs to calibrate governance, policies, procedures and life cycle controls to the size and nature of their activities, their use of AI and their risk profile. In assessing the potential impact, the FI should consider the potential financial, operational, regulatory, legal or reputational impact on the FI, and the potential impact on customers or other stakeholders, including in relation to fairness, ethical conduct and consumer protection.

Unless the basic governance track applies, a FI should establish frameworks for AI oversight, identification, inventory and materiality assessment, apply relevant life cycle controls and maintain adequate capabilities and technology infrastructure.

A FI may use basic AI governance policies where poor performance or unavailability of the AI service or tool is unlikely to have a material adverse impact on the FI, its customers or other stakeholders. Examples include using AI to draft or rephrase customer emails, summarise documents or

meeting notes for internal reference, assist with initial document review, generate internal formulas or visualisations, design marketing or internal materials, or locate internal policies and procedures.

Basic AI governance policies and procedures should minimally include: (a) clear accountability for AI oversight; (b) permitted and prohibited uses, including effective human oversight and any mandatory review of outputs; (c) an approved list of AI tools and a process for approving new tools; (d) staff education; (e) regular compliance checks; and (f) periodic and trigger-based review of whether the use remains low risk.

Basic AI governance policies and procedures are therefore not an exemption from governance. It remains important that personnel review and check outputs, that confidential, proprietary or client information is not shared with public AI tools where prohibited, and that the FI reassesses the use if its impact or risk profile changes.

2. Key supervisory expectations

2.1 Board and senior management oversight

The board and senior management should ensure that the FI's existing risk management frameworks, policies and practices identify, assess and address risks posed by AI. They should maintain effective oversight, foster an appropriate risk culture and ensure that AI use does not conflict with regulatory requirements or other supervisory expectations.

The board, or a committee delegated by it, should approve and regularly review the overall governance approach, strategic direction and key frameworks for AI risk management. Risks from AI use should be addressed explicitly within the FI's risk appetite framework, including qualitative statements and quantitative measures where appropriate, such as thresholds for AI-related incidents, provider concentration or breaches of performance thresholds.

The board and senior management should also set clear roles and responsibilities, maintain an adequate understanding of AI to provide oversight and effective challenge, and regularly review the FI's approach, risk appetite, capabilities and culture as AI developments, business strategies and the FI's risk profile change.

Senior management is responsible for ensuring effective implementation of the framework, instituting controls across the AI life cycle, establishing roles across business and control functions, maintaining escalation processes for material risks and exceptions, updating the board on material issues, and ensuring adequate personnel, technology, financial resources, training and capacity building.

Where group frameworks are used, local senior management remains accountable for ensuring that those frameworks address Singapore-specific considerations and regulatory requirements. Local management should have visibility into and input on Singapore matters, clear escalation paths and the ability to demonstrate to MAS how its oversight responsibilities are discharged.

2.2 AI identification, inventory and risk materiality assessment

FIs should identify AI use consistently across relevant business and functional areas, including AI incorporated within or used in the delivery of material third-party services. A designated control function should retain independent oversight of the process and be the final arbiter of whether a given use constitutes AI use.

The FI should maintain an accurate AI inventory, linked where relevant to other inventories such as data and third-party registers. Attributes may include purpose and description, approved scope of use, model type, data used, dependencies, life cycle status, risk materiality rating, model review status, key roles and responsibilities, and links to supporting documentation.

The Guidelines also address residual risks arising where AI cannot be consistently identified, including where AI features are embedded in third-party services or because of shadow AI. Those residual risks should remain within the FI's risk appetite, supported by mitigating measures such as staff guidance, restrictions on public AI tools, network monitoring or data loss prevention controls.

Risk materiality should be assessed consistently for each AI use case in the FI, taking into account inherent risk and residual risk before and after risk management controls are applied. At a minimum, the methodology should consider impact, complexity and reliance, including the level of autonomy and human involvement. Residual risk should be within the FI's risk appetite before deployment.

2.3 AI life cycle controls

FIs should plan for and implement controls across the entire AI life cycle, from inception through retirement or decommissioning, with clear responsibilities across different business and functional units. Controls may be calibrated to the assessed risk materiality, but high-risk use cases should have contingency plans with fallback systems or manual processes and kill switches should have clear activation protocols implemented which are subject to regular testing.

Data management controls should ensure data used across the AI life cycle is fit for purpose and representative, of high quality, and subject to robust data governance. General data governance standards should be

applied and, where necessary, uplifted to address the variety and volume of data used across the AI life cycle.

Transparency and explainability should be tailored to the context and nature of the AI use case and its assessed risk materiality. FIs should define fair outcomes and address harmful bias and discriminatory outcomes, with greater attention to areas such as credit decisioning and insurance underwriting which pose higher risks of unfair access or denial of financial services to individuals. Human oversight should include clear responsibilities, competent personnel with authority to intervene, escalation mechanisms and records of decisions and interventions.

Selection, evaluation and testing of AI systems should be aligned with the use case and risk materiality. FIs should set measurable thresholds, test across plausible conditions and edge cases, address failure modes and guardrails for Generative AI and AI agents, and consider simpler or better-understood alternatives where appropriate.

Technology and cybersecurity controls should support secure deployment, access control, network segmentation, input validation, authentication, encryption, data loss prevention and appropriate controls over third-party components. The AI development process should be sufficiently documented to support reproducibility and auditability, including data sources, code and environments, testing results, explainability and fairness analysis, assumptions, limitations and mitigants.

Before deployment, the AI use case should be reviewed by parties not involved in its development. The scope and independence of review should be proportionate to risk materiality, with formal independent validation required before deployment for high-risk use cases. Technology and cybersecurity reviews should also confirm controlled and secure production readiness.

After deployment, FIs should continue to monitor performance, robustness, stability, data quality and fairness, with thresholds and checks for data, concept and model drift. They should maintain incident processes, records and feedback channels, and consider kill switches or override mechanisms for high-risk use cases. High-risk use cases should undergo regular independent re-validation or review where necessary. For AI use cases not assessed as high risk, other forms of documented review, such as peer reviews by qualified individuals not involved in development or deployment, may be undertaken.

Robust controls should be implemented for managing changes to deployed AI. The controls should assess the significance of changes, require review and re-approval for significant changes, prevent unauthorised alterations and support traceability. For AI systems designed to be updated automatically, there should be clear definitions of what can be updated automatically. When AI is retired or decommissioned, the FI should

address dependencies, data retention, secure removal and appropriate stakeholder notifications. Where changes may be implemented without prior review, compensating controls such as enhanced monitoring, should be applied.

Taken together, these requirements indicate that life cycle controls should be documented as an integrated control framework rather than treated as isolated model-development checks. The degree of formality, independence, testing and monitoring should reflect the potential impact, complexity and reliance associated with the relevant use case.

The table below summarises relevant AI life cycle controls and the way in which the Guidelines contemplate their calibration

Area	Key expectations	Calibration
Data management	Fit-for-purpose, representative and high-quality data; appropriate classification, security, privacy, sourcing, processing and lineage.	Apply existing data standards and uplift them where necessary for the variety, volume and sensitivity of AI data.
Transparency, fairness and human oversight	Provide appropriate explanations; define fair outcomes; identify and mitigate bias; assign competent personnel with authority to intervene.	Increase transparency, explainability and human review for credit, underwriting, advisory and other higher-impact uses.
Selection, evaluation, testing and security	Select systems and models against use-case objectives; test reliability, safety, edge cases and failure modes; deploy in secure environments with access controls and guardrails.	Use measurable thresholds and more rigorous evaluation for complex, novel or higher-risk AI, including Generative AI and AI agents.
Pre-deployment review, monitoring and change	Independent review before deployment; monitor performance, drift, incidents and fairness; control changes, re-validate and maintain escalation and override processes.	High-risk use cases require formal independent validation before deployment and regular independent re-validation or review where necessary.

Area	Key expectations	Calibration
Reproducibility, contingency and retirement	Maintain documentation sufficient for auditability and replication; test fallback options; use kill switches where appropriate; address dependencies, data retention and secure removal.	Increase documentation, testing and contingency planning where failure could materially affect the FI, customers or other stakeholders.

The table should be read alongside the requirement for clear allocation of responsibilities. Business and functional units may carry out particular control activities, but the FI should retain appropriate independent oversight and ensure that controls remain effective throughout the life cycle.

In this respect, life cycle controls are not limited to the underlying model. They extend to the data, system architecture, interfaces, third-party components, human processes and operational environment through which the AI use case is delivered.

FIs should therefore map the relevant controls to the actual operating model, rather than assume that a policy-level statement will address risks arising from deployment, use or change.

2.4 Third-party AI

The decision to onboard and use third-party AI ultimately rests with the FI, and the FI retains primary accountability for its use, its impact on stakeholders and its ability to comply with regulatory requirements. MAS's outsourcing and third-party service expectations also apply to third-party AI.

Contracts should give the FI a risk-appropriate degree of visibility over the introduction of AI and subsequent updates or changes. FIs should test third-party AI in the context of their own use cases, including with their own data where appropriate, and maintain documentation of how the service was tested and deemed suitable.

Where provider disclosures are inadequate, compensatory testing or greater human oversight may be needed. Independent certifications or external assessments may assist only when performed by independent and competent parties, and self-attestations will not provide adequate coverage of key risks from AI use.

FIs should also assess supply-chain risks, model and data provenance, concentration risk, change notifications, contingency arrangements and the

complexity of the service. Legal agreements should address matters such as performance guarantees, data protection, audit rights, notification when AI is introduced or updated, and the FI's consent before AI is incorporated, where appropriate.

The FI should ensure that the residual risk materiality of the use case which uses or depends on third-party AI is within the FI's risk appetite. Where this cannot be achieved, the FI should consider limiting or suspending the third-party AI service, or replacing the relevant third-party provider.

2.5 Capability and capacity

FIs should ensure that personnel involved in developing, deploying and using AI have the necessary competence, proper conduct, training and support. Capability and capacity programmes should be reviewed regularly and updated for newer AI technologies, with resources proportionate to the risk profile of the relevant use case.

Technology infrastructure should be adequate for the intended AI use, including sufficient hardware, software, network infrastructure, system memory and secure data pipelines to meet performance, scalability, resilience, safety and cybersecurity needs. FIs should also take account of applicable technology risk guidance and relevant industry frameworks.

3. Relationship with existing frameworks

The Guidelines are intended to complement, rather than displace, existing financial-sector requirements. FEAT principles continue to apply, and FIs should continue to comply with the Fair Dealing Guidelines when using AI to deliver products and services.

The Guidelines also sit alongside MAS's Technology Risk Management Guidelines and requirements relating to outsourcing and third-party services. Data use should be considered in conjunction with the Personal Data Protection Commission's advisory guidelines on personal data in AI recommendation and decision systems and in Generative AI.

FIs may also refer to the IMDA Model AI Governance Framework, including its work on agentic AI, the Cyber Security Agency of Singapore's guidelines and companion guide on securing AI systems, IMDA's Starter Kit for Testing LLM-Based Applications for Safety and Reliability, and the NIST AI Risk Management Framework.

Global group frameworks may be leveraged where they meet the Guidelines' expectations. The practical question is therefore not whether a FI must build a wholly separate framework, but whether its existing

arrangements provide sufficient local accountability, visibility, documentation, control and escalation.

4. Implications for FIs

For FIs, the immediate implication is likely to be a governance uplift. FIs will need a comprehensive view of where AI is used, who is accountable, how materiality is assessed and how controls operate from selection through retirement. The designated control functions should have sufficient independent oversight to arbitrate risk materiality assessments and to ensure such assessments are regularly reviewed for accuracy and relevance.

AI used in credit scoring, insurance underwriting, financial advisory or fund management is likely to require greater attention to explainability, fairness, human oversight and redress because of its potential impact on customer outcomes. The calibration should be tied to the actual use case and not simply to the label attached to the technology.

Generative AI and AI agents may require additional controls because of their potential for hallucination, data leakage, adversarial attack, unauthorised actions and over-reliance. FIs should consider whether existing human oversight, incident management, testing and cybersecurity arrangements remain adequate for these technologies.

The resulting work programme will differ across FIs. A large institution with material, customer-facing or autonomous use cases may require more formal inventories, validation, monitoring and contingency arrangements, while lower-impact internal assistance may be managed under the basic governance track.

Conclusion

The Guidelines provide a structured supervisory framework for the use of AI by FIs. They make explicit the governance, identification, materiality, life cycle and capability disciplines that MAS expects FIs to apply to AI, while allowing calibration according to the FI's size, activities and risk profile.

The Guidelines therefore require more than a review of model risk policies. FIs should be able to demonstrate where AI is used, how materiality is determined, who is accountable, what controls operate through the life cycle and how the FI will respond to failure, change or new risks.

FIs should therefore treat the period before 7 October 2027 as an opportunity to establish the governance and identification foundations, while developing more extensive life cycle controls, capabilities and infrastructure required by 7 October 2028. For FIs, the central task will be to convert the supervisory expectations into a proportionate and demonstrable risk management framework.

How Drew & Napier LLC can assist

Drew & Napier's Financial Services Regulatory practice advises FIs and other market participants on AI governance, technology risk and Singapore financial regulatory matters.

Our work includes advising on AI governance frameworks, risk appetite and accountability, AI identification and inventories, life cycle controls, third-party and embedded AI, technology and operational risk, data protection, vendor contracting and regulatory engagement with MAS.

We can assist FIs in assessing the implications of the Guidelines, carrying out gap analyses, developing proportionate governance and control frameworks, reviewing vendor arrangements, preparing implementation plans and engaging with MAS on regulatory questions.

The content of this article does not constitute legal advice and should not be relied on as such. Specific advice should be sought about your specific circumstances. Copyright in this publication is owned by Drew & Napier LLC. This publication may not be reproduced or transmitted in any form or by any means, in whole or in part, without prior written approval.

Special thanks to Kyra Tan, Legal Intern, for her diligent assistance with this publication.

Please feel free to reach out to us if you have any query in this area:

Grace Chong
Head, Financial Services Regulation
Director, Corporate & Finance
T: +65 8869 0445
E: grace.chong@drewnapier.com

Bryan Ong
Senior Associate, Corporate & Finance
T: +65 8058 8905
E: bryan.ong@drewnapier.com

Shireen Lee
Trainee, Litigation
E: shireen.lee@drewnapier.com

Drew & Napier LLC

10 Collyer Quay
#10-01 Ocean Financial Centre
Singapore 049315

www.drewnapier.com

T: +65 6535 0733
T: +65 9726 0573 (After Hours)
E: mail@drewnapier.com

DREW & NAPIER