

Implementing the Monetary Authority of Singapore's AI Risk Management Guidelines



17 October 2026

On 7 October 2026, the Monetary Authority of Singapore (the “**Authority**”) issued its [Guidelines on Artificial Intelligence Risk Management](#) (the “**Guidelines**”) for financial institutions (“**FIs**”) using artificial intelligence (“**AI**”).¹ The Guidelines inform the Authority’s [supervisory assessment](#) of an FI, although they are not a legally binding notice.² An FI should implement the Guidelines in a manner commensurate with the size and nature of its activities, and the extent to which its use of AI could lead to material risks to it.³ In this article, we discuss the practical decisions FIs need to make when implementing the Guidelines.

The Guidelines’ definition of AI covers systems and models that produce outputs using what they have learned from data, including generative AI and agents.⁴ Tools whose outputs depend solely on predefined programming logic or rules, such as calculators, are excluded.⁵

Under the Guidelines, AI risk materiality is assessed by use case, including how the FI relies on AI and the extent of human involvement.⁶ For example, an FI may approve a tool to help employees find policies, provided employees check the original documents before acting. Allowing the tool to answer customers without staff review changes both the task and the protection against errors. Testing that shows the tool retrieves the right policy does not show it gives customers complete and accurate answers.

¹ Guidelines, paragraph 1.1.

² Authority, Supervisory Approach and Regulatory Instruments, Regulatory Instruments, items 3(b) and 4; Guidelines, paragraph 1.1.

³ Guidelines, paragraph 1.6.

⁴ Guidelines, paragraph 1.3; footnote 6.

⁵ Guidelines, paragraph 1.3.

⁶ [Response to Feedback](#), paragraphs 7.5–7.6, 10.30; Guidelines, paragraphs 4.10–4.12(c).

Executive Summary

- **The Guidelines apply from 7 October 2027** for basic governance, oversight, inventory and risk materiality; lifecycle controls and capability follow by 7 October 2028 where applicable.⁷ The Authority expects appropriate lifecycle controls for AI use cases assessed as having high risk materiality as soon as possible.⁸
- **Basic governance is available only where poor performance or unavailability of the AI services or tools used by the FI is unlikely to materially harm the FI, its customers or other stakeholders.**⁹ Otherwise, the FI implements the oversight and risk management expectations in Sections 3–6, with lifecycle controls calibrated to the risk materiality of each use case.¹⁰
- **Approval should define the task and permitted reliance on AI.**¹¹ Testing should support that use; use cases with high inherent risk materiality need formal independent validation before deployment.¹²
- **Relying on another group entity to maintain the FI’s AI inventory or perform its AI risk assessments may constitute outsourcing.**¹³ Third-party AI controls can also apply where no service is outsourced.¹⁴

Who Are the Guidelines For?

The Guidelines adopt the FI definition in the [Financial Services and Markets Act 2022](#), which includes banks, insurers, capital markets services licensees, financial advisers and regulated payment providers, including licensed digital payment token service providers.¹⁵ For locally incorporated FIs under consolidated supervision or owning critical information infrastructure, the Guidelines also apply on a group basis to their branches and subsidiaries.¹⁶

When May an FI Apply Basic Governance?

Paragraph 2.3 allows basic governance where poor performance or unavailability of the AI services or tools used by the FI is unlikely to have a material adverse impact on the FI, its customers or other stakeholders.¹⁷

Paragraph 2.5(f) calls for periodic and trigger-based review of “whether the FI’s AI use continues to meet the criteria in paragraph 2.3.”¹⁸ An FI using a drafting copilot and also operating a customer-facing chatbot assessed as outside paragraph 2.3’s low-impact criterion should apply Sections 3–6 to both uses.¹⁹ The

⁷ Guidelines, paragraph 1.8.

⁸ Response to Feedback, paragraphs 13.5–13.6.

⁹ Guidelines, paragraphs 2.2–2.3, 2.5(f).

¹⁰ Guidelines, paragraph 2.2(b) and footnote 12; Response to Feedback, paragraphs 3.6–3.10.

¹¹ Guidelines, paragraphs 4.7, 5.21.

¹² Guidelines, paragraphs 5.14–5.15, 5.19; Response to Feedback, paragraph 7.21.

¹³ [Outsourcing Guidelines](#), paragraph 3.1 and Annex 2.

¹⁴ Guidelines, paragraphs 5.10–5.11.

¹⁵ Guidelines, footnote 1.

¹⁶ Guidelines, paragraph 1.2; Response to Feedback, paragraph 2.3.

¹⁷ Guidelines, paragraphs 2.2–2.3, 2.5(f); Response to Feedback, paragraphs 3.7–3.10.

¹⁸ Guidelines, paragraph 2.5(f).

¹⁹ Guidelines, paragraph 2.2; Response to Feedback, paragraphs 3.7–3.10.

copilot can receive lighter testing and review, while data management, safety and cybersecurity controls still apply proportionately.²⁰

Paragraph 2.4 of the Guidelines gives examples of AI uses that would generally meet the criterion for basic governance.²¹ These include helping staff draft customer emails, summarise meeting notes or initially analyse documents for internal reference, with humans checking the outputs before use.²²

Where human checking supports the low-impact assessment, the FI can sample completed reviews to establish whether staff identify errors and intervene.²³ A sign-off alone does not establish whether staff checked the source or accepted a plausible answer, which affects the basis of the low-impact assessment. The Guidelines recognise that automation bias and decision fatigue can weaken human oversight as AI use increases.²⁴

Basic governance should cover accountability, permitted and prohibited uses (including restrictions on confidential, proprietary or client information in public AI tools), human checking, approved tools and an approval process, staff education, compliance checks and reassessment.²⁵ The Authority encourages at least annual review as good practice; the FI should document a review that identifies no necessary updates.²⁶

An FI that has not adopted AI is still expected to establish basic policies on whether and how staff may use AI, to address unauthorised or “shadow” AI use.²⁷ Where Section 4 applies, the FI should identify and mitigate risks arising from unidentified AI use and keep the residual risk within its risk appetite.²⁸

The rest of this article focuses on FIs that cannot rely on basic governance alone.

What Does the FI Need to Approve?

The FI’s approval should identify the permitted use and any conditions on reliance.²⁹ For an agent, we recommend that the FI record which systems and tools the agent may access and what actions the agent may take.³⁰

Before approval, the FI should identify its AI use, including AI embedded in services from material third-party providers.³¹ Its inventory should include low-impact uses to the extent possible or practicable, with detail proportionate to the nature and risk materiality of the use.³²

²⁰ Guidelines, paragraph 2.2(b) and footnote 12; Response to Feedback, paragraphs 3.9–3.10.

²¹ Guidelines, paragraph 2.4.

²² Guidelines, paragraph 2.4 and footnote 13.

²³ Guidelines, paragraphs 1.10(f), 5.9(d).

²⁴ Guidelines, paragraph 5.9.

²⁵ Guidelines, paragraph 2.5.

²⁶ Response to Feedback, paragraph 3.16(c).

²⁷ Response to Feedback, paragraph 3.14.

²⁸ Guidelines, paragraph 4.4.

²⁹ Guidelines, paragraphs 4.7, 5.21.

³⁰ Guidelines, paragraph 4.7; footnote 25; Response to Feedback, paragraph 7.7.

³¹ Guidelines, paragraph 4.2.

³² Guidelines, paragraphs 2.2(a), 4.2, 4.5, 4.7; Response to Feedback, paragraphs 6.15–6.16.

The FI should assess inherent risk before risk management controls and residual risk after controls, considering impact, complexity and reliance as minimum dimensions.³³ An outsourcing rating does not answer that question: AI use-case materiality, material outsourcing and material third-party providers are different assessments.³⁴

Human involvement informs the assessment, but does not automatically make a use low risk.³⁵ Where inherent risk materiality is high, the specific expectations for independent validation and contingency planning apply even if controls reduce residual risk to an acceptable level.³⁶

The FI's testing should establish whether the AI performs the proposed task on representative data, including difficult conditions and known exceptions.³⁷ For agents, this includes testing failure modes and whether guardrails prevent actions outside the intended limits.³⁸ The residual risk should be within the FI's risk appetite before deployment.³⁹

Who Approves and Who Challenges?

The board or its delegated committee approves and regularly reviews the AI governance approach and ensures that AI risks are addressed in the FI's risk appetite framework.⁴⁰ Senior management or relevant committees may approve the materiality methodology and lifecycle controls; the FI may place approval of AI use cases assessed as having high risk materiality with the board or senior management.⁴¹ Senior management assigns responsibility for AI identification, inventory and materiality assessment, establishes escalation processes for material risks and exceptions, and updates the board on material AI risk issues in a timely manner.⁴²

If the same team develops, configures and operates a tool, the FI needs to decide who can independently challenge the proposed use. Reassigning the review within that team may leave the independence problem unresolved. The challenge should address whether the testing supports the intended reliance, including the protection attributed to human review.

For AI use cases assessed as having high risk materiality, the Guidelines call for formal validation before deployment by suitably qualified internal or external personnel independent of development and deployment.⁴³ Other uses need pre-deployment review, which may include proportionate, documented peer review by people who did not develop, deploy or operate the AI.⁴⁴

³³ Guidelines, paragraphs 4.10–4.12.

³⁴ Guidelines, paragraphs 4.2, 4.10; Outsourcing Guidelines, paragraph 3.1; Response to Feedback, paragraphs 9.19–9.21.

³⁵ Guidelines, paragraph 4.12(c); Response to Feedback, paragraphs 7.5–7.6, 7.19.

³⁶ Guidelines, paragraphs 5.3, 5.19; Response to Feedback, paragraphs 7.20–7.21.

³⁷ Guidelines, paragraphs 5.14–5.15.

³⁸ Guidelines, paragraph 5.15; footnote 46.

³⁹ Guidelines, paragraph 4.11.

⁴⁰ Guidelines, paragraphs 3.4(a)–(b).

⁴¹ Response to Feedback, paragraphs 4.8, 4.10(c).

⁴² Guidelines, paragraphs 3.5(d)–(f).

⁴³ Guidelines, paragraphs 5.18–5.19; Response to Feedback, paragraph 10.55.

⁴⁴ Guidelines, paragraphs 5.18–5.20; footnote 48.

Where the FI uses a three-lines model, internal audit can also provide independent assurance over the effectiveness of the AI risk management framework.⁴⁵

What If the FI Relies on Its Group?

An FI may use a group framework that meets the Guidelines' expectations, with oversight from an appropriate regional or global committee.⁴⁶ A group-level control function is also permitted where the FI demonstrates adequate oversight of AI used in Singapore.⁴⁷ Global pre-deployment reviews can be relied upon where they meet the relevant expectations, including formal independent validation for AI use cases assessed as having high risk materiality and testing addressing the local context; gaps require additional local review or testing before deployment.⁴⁸ Local senior management should be able to demonstrate how it discharges its oversight responsibilities when relying on group frameworks.⁴⁹

Group reliance may also require an outsourcing assessment. The Outsourcing Guidelines include group entities as service providers and identify risk management and compliance as examples of outsourced services.⁵⁰ Where the Singapore FI depends on another group entity to perform those functions on an ongoing basis, the arrangement may constitute outsourcing if the service is integral to the FI's financial services.⁵¹ A discrete independent assessment may fall outside that definition.⁵² Where a regional team performs the Singapore FI's risk assessments, the FI should identify how it will monitor the service and challenge the conclusions.

The Outsourcing Guidelines for non-bank FIs allow proportionate intragroup due diligence, but retain institution-specific expectations for monitoring, continuity, audit, access and local compliance; banks and merchant banks have a separate [bank framework](#).⁵³

The Authority has proposed replacing the Outsourcing Guidelines with [Guidelines on Third-Party Risk Management](#); FIs should revisit their assessments when the final framework is issued.⁵⁴

What Information Is Needed from Providers?

Third-party AI controls apply to externally developed models and services incorporating AI, including models hosted within the FI, even where no service is outsourced.⁵⁵

The FI can start by asking each relevant third-party provider for a list of the services supplied, where those services use AI and the assurance information available.

⁴⁵ Guidelines, footnote 19.

⁴⁶ Guidelines, paragraphs 1.5, 3.6; footnote 17.

⁴⁷ Response to Feedback, paragraph 8.4.

⁴⁸ Response to Feedback, paragraph 10.59.

⁴⁹ Guidelines, paragraph 3.6.

⁵⁰ Outsourcing Guidelines, paragraph 3.1; Annex 2, paragraphs 1(c) and 1(m).

⁵¹ Outsourcing Guidelines, paragraph 3.1.

⁵² Outsourcing Guidelines, Annex 2, paragraph 2(c).

⁵³ Outsourcing Guidelines, paragraphs 5.4.5, 5.11.

⁵⁴ Response to Feedback, footnote 3, page 31; Consultation Paper on Proposed Guidelines on Third-Party Risk Management, paragraphs 1.2–1.3.

⁵⁵ Guidelines, footnote 10; paragraphs 5.10–5.11.

Consistent vendor and service identifiers can link the AI inventory to existing outsourcing and third-party registers.⁵⁶ The FI should distinguish independent assessments from a provider's self-attestations, and assess performance on its own use and data.⁵⁷ Open-source dependencies and concentration on a single provider also need assessment.⁵⁸

A provider-commissioned review may qualify as documented review, but counts as formal independent validation only where the reviewer has the necessary expertise and independence from the provider's development and deployment teams.⁵⁹

For example, a provider checking customer files may not disclose its training and testing details. The FI can test records with known omissions and records known to be complete to measure missed exceptions and false alarms. If only flagged files receive human review, missed exceptions in unflagged files may remain unexamined. We recommend testing samples of unflagged files against independently established outcomes. Testing cannot close every disclosure gap, and residual risk must remain within the FI's risk appetite.⁶⁰

Contracts should give the FI enough information about AI introduced into the service and later updates to assess the resulting risks.⁶¹ Where appropriate, the FI could seek a term requiring the provider to obtain the FI's agreement before introducing AI into the service.⁶² Where the service is an outsourcing arrangement, the FI also needs the applicable register entry and provisions on confidentiality, security, subcontractors, business continuity, audit and access.⁶³

If residual risk remains unacceptable, the FI may need to restrict or suspend use, or replace the tool or provider.⁶⁴

When Does the Assessment Need to Change?

Reliance can change without a software update: staff may stop checking sources, or an agent may gain permission to change customer records. The FI needs to address failures in required human checking. Changes to permitted reliance or agent permissions can alter the scope of use and should be assessed for significance under paragraph 5.25(a).⁶⁵

The FI should review and reapprove significant changes before implementation, and set enhanced controls and boundaries for automatic updates.⁶⁶ The FI's monitoring should examine oversight decisions and interventions, including incidents and near misses, as well as relevant actions and tools used by agents.⁶⁷ For

⁵⁶ Guidelines, footnote 23.

⁵⁷ Guidelines, paragraphs 5.10–5.11(a).

⁵⁸ Guidelines, paragraphs 5.11(b)–(c).

⁵⁹ Response to Feedback, paragraph 10.58.

⁶⁰ Guidelines, paragraphs 4.11, 5.2, 5.10.

⁶¹ Guidelines, paragraph 5.10.

⁶² Guidelines, paragraph 5.11(f).

⁶³ Outsourcing Guidelines, paragraphs 4.1.1, 5.5–5.9; Annex 4, paragraph 1.

⁶⁴ Guidelines, paragraphs 5.2 and 5.11 (opening paragraph).

⁶⁵ Guidelines, paragraph 5.25(a); footnote 49.

⁶⁶ Guidelines, paragraphs 5.24–5.25.

⁶⁷ Guidelines, paragraphs 5.9(d), 5.23(a).

AI use cases assessed as having high risk materiality, the FI should prepare and test contingency plans, consider kill switches or overrides, and test the activation protocols where kill switches are used.⁶⁸

AI use cases assessed as having high risk materiality should undergo regular re-validation by parties independent of development and deployment.⁶⁹

Limited pilots may use adjusted controls, such as time or user limits, success criteria and close monitoring; AI use cases assessed as having high risk materiality still need formal independent validation before deployment.⁷⁰

Existing [FEAT](#), [Fair Dealing](#), [technology risk management](#), data-protection and [GenAI cyber-risk guidance](#) continue to apply.⁷¹ An AI-related incident may also trigger existing notification duties under the FI's applicable [Technology Risk Management Notice](#) or business continuity requirements; the Guidelines introduce no separate AI-specific reporting regime.⁷²

Implementation Resources

The [MindForge AI Risk Management Toolkit](#), announced by the Authority on 20 March 2026, includes an [Operationalisation Handbook](#) and [Implementation Examples](#) with templates and case studies.⁷³ The December 2024 [Artificial Intelligence Model Risk Management](#) information paper provides further good practices.⁷⁴ Because the Handbook was prepared against the proposed Guidelines, an FI using its mapping should check the references against the final text.⁷⁵

⁶⁸ Guidelines, paragraphs 5.3, 5.23(b); footnote 29.

⁶⁹ Guidelines, paragraph 5.24; Response to Feedback, paragraph 10.68.

⁷⁰ Guidelines, footnotes 28, 45; Response to Feedback, paragraphs 10.77–10.79.

⁷¹ Guidelines, paragraphs 1.2, 5.4, 5.16; footnotes 2, 9, 15, 30, 47.

⁷² Technology Risk Management Notices, including Notices [FSM-N13](#) and [FSM-N21](#), paragraphs 1–2 and 7; Response to Feedback, paragraph 10.69.

⁷³ MindForge announcement, 20 March 2026; Handbook, appendices E–H; Implementation Examples.

⁷⁴ Artificial Intelligence Model Risk Management, 5 December 2024.

⁷⁵ Handbook, introduction.



What Should an FI Prepare Now?

We suggest that FIs prepare the following six records:

1. **An AI inventory.** Reconcile internal uses, embedded AI and provider information with existing technology and outsourcing records.
2. **A governance assessment.** Record the paragraph 2.3 assessment, the materiality methodology, controls for each use and, where Section 3 applies, board or delegated committee approval of the governance approach.
3. **A use and testing record.** Define the task, users, data, permitted actions and any human checking relied upon, with acceptance thresholds and test results.
4. **A group and provider assessment.** Record services performed for the FI, outsourcing classifications, evidence gaps and contractual restrictions.
5. **An approval record.** Record the approving body, permitted reliance, conditions on use and completed independent review or validation.
6. **A monitoring plan.** Identify who can intervene, monitoring and re-validation schedules, escalation or retesting triggers, and responses to errors, outages or changes in use.

Practical Implementation Support

HM Strategy can help FIs implement the Guidelines by defining the permitted use of AI, assessing the evidence supporting reliance, and developing the resulting governance, testing and monitoring arrangements. This can include AI inventories, risk materiality assessments, group and provider assessments, and approval records. Subject to appropriate independence arrangements, we can also provide independent assurance over the design and operation of AI governance through our outsourced internal audit services.

For further information, contact Chris Holland at chris.holland@hmstrategy.com.

Disclaimer: The material in this post represents general information only and should not be relied upon as legal advice. Holland & Marie Pte. Ltd. is not a law firm and may not act as an advocate or solicitor for purposes of the Singapore Legal Profession Act.