



Monetary Authority of Singapore

Response to Feedback Received

P017-2025 – 7 October 2026

Response to Feedback Received on Guidelines on Artificial Intelligence Risk Management



Contents

1. Preface	3
2. Applicability of the Guidelines	4
3. Proportionality	6
4. Responsibilities of the Board and Senior Management	10
5. Establishing a dedicated cross-functional committee	13
6. AI Identification & Inventory	15
7. Risk materiality assessment	20
8. Roles and responsibilities for AI risk management systems, policies and procedures	25
9. Third-party AI management	27
10. Other lifecycle controls	31
11. AI Capability and Capacity	46
12. Generative AI and AI Agents	49
13. Implementation Timeline	51
Annex A	52



1. Preface

- 1.1. On 13 November 2025, MAS issued a consultation paper on the proposed Guidelines on Artificial Intelligence (AI) Risk Management (the “Guidelines”) to set out its supervisory expectations relating to AI risk management in the financial sector. The Guidelines, which apply to all financial institutions (FIs), focus on oversight of AI risk management, key AI risk management systems, policies and procedures, key AI life cycle controls, as well as capabilities and capacity needed for the use of AI.
- 1.2. The consultation period closed on 31 January 2026, and MAS would like to thank all respondents for their contributions. The list of respondents is in Annex A.
- 1.3. MAS has carefully considered the feedback received and appreciates the views shared. Where appropriate, the feedback has been reflected in the final Guidelines, having regard to the Guidelines’ overall objectives. The final Guidelines can be accessed via the following link: *[Guidelines on Artificial Intelligence Risk Management](#)*.
- 1.4. Comments that are of wider interest, together with MAS’ responses are set out below.



2. Applicability of the Guidelines

Scope of Guidelines

- 2.1. MAS had proposed to apply the Guidelines to all FIs. FIs in Singapore that are branches or subsidiaries with parent entities in other jurisdictions may leverage the AI risk management frameworks of their parent entities, as long as such frameworks meet the expectations set out in the Guidelines.
- 2.2. There was broad support for the proposal. Two respondents sought clarification on whether the Guidelines would apply to FIs which are incorporated outside of Singapore and foreign subsidiaries of FIs incorporated in Singapore.

MAS' Response

- 2.3. The expectations on AI risk management set out in the Guidelines are relevant to all FIs deploying AI, and apply regardless of where an FI is incorporated. The Guidelines apply on a group basis for an FI with a branch or subsidiary under it, and which is: i) subject to consolidated supervision by MAS¹; or ii) an owner of critical information infrastructure².

Scope of AI

- 2.4. MAS proposed the scope of AI to cover AI use cases, systems and models. This extends to third-party AI, which encompasses all providers of third-party AI services, covering third-party systems, models, as well as data used for AI. Existing third-party services where AI has been introduced would also fall within scope.
- 2.5. Some respondents sought clarification on whether certain use cases, systems and models fall within the proposed scope of AI. In particular, respondents asked whether hybrid systems with both rule-based and AI components, embedded AI incorporated within, or used in the delivery of, third-party services not explicitly marketed or sold as AI, conventional risk models (e.g. logistic regression), and AI in pilot or proof-of-concept stage, are within scope, and whether the scope of AI is limited to only use

¹ These include locally-incorporated FIs from the banking and insurance sectors.

² As defined in the Cybersecurity Act 2018.

cases with customer impact. A few respondents also requested examples to illustrate what is excluded from the scope of AI, such as stochastic models (e.g. Monte Carlo simulations) or rule-based tools like Robotic Process Automation (RPA).

- 2.6. On the scope of third-party AI, a few respondents requested to limit it to third-party procured AI solutions, with either lighter requirements or exclusions to be applied for other forms of third-party AI, such as embedded AI.

MAS' Response

- 2.7. For the purpose of the Guidelines, AI refers to machine-based systems or models that derive outputs through learned premises such as the data or inputs it receives. The outputs can include estimates, predictions, content, summaries, recommendations, or decisions that may influence physical or virtual environments. Different AI vary in their levels of autonomy and adaptiveness after deployment. Calculators or tools whose outputs are solely based on predefined programming logic or rules would not be regarded as AI for the purpose of the Guidelines.
- 2.8. MAS has provided illustrative examples of what would generally fall within scope of the AI definition below. Such examples are meant to be illustrative and are neither prescriptive nor exhaustive, and newer techniques and applications not listed here should be assessed against this definition.

Illustrative examples of what would generally fall within scope of the AI definition	Illustrative examples of what would generally fall outside scope of the AI definition
• machine learning (including supervised, unsupervised, and reinforcement learning) such as logistic regression and gradient boosting; • deep learning; • natural language processing; • computer vision; • Generative AI (including large language models); and • AI agents (including multi-agent systems).	• standard rule-based investment formulas and portfolio construction methods; • expert systems based on explicit if-then rules; • traditional robotic process automation (RPA) following rule-based decision trees; and • stochastic simulation techniques such as Monte Carlo simulations that are not derived from training data.

- 2.9. MAS' intent is to ensure that risks from AI use that may materially affect an FI's operations, decision-making, or service delivery are appropriately identified and managed, regardless of whether the AI is developed in-house or provided by a third party. Against this context, the scope of "AI" under



the Guidelines covers all use cases, systems, and models that incorporate any element of AI. This includes hybrid systems, and AI that is partially deployed (i.e. in pilot or proof-of-concept stage).

- 2.10. With regard to third-party AI, risks may arise not only from AI services procured from third parties, but also where AI is incorporated within third-party services, or used by third-party service providers in delivering services to the FI. For example, an FI may rely on AI features incorporated within software-as-a-service to extract information for key business decisions, or a third-party service provider may use AI to automate key processes which affect the reliability of services provided to the FI.
- 2.11. Accordingly, MAS' intent is that the scope of third-party AI should include AI incorporated within, or used in the delivery of, services to the FI. This reflects the need for an FI to consider these risks given its implications on the FI. MAS recognises that such AI may not always be readily visible to an FI, particularly where the AI is not explicitly marketed or provided as an AI service, and will allow an FI to adopt a risk-proportionate approach towards the identification of third-party AI. This is further discussed under Section 6 of this response paper.

3. Proportionality

- 3.1. MAS had proposed that an FI should implement the Guidelines in a risk-based and proportionate manner. All FIs should establish basic AI governance policies and procedures for the use of AI commensurate with its level of AI adoption. The full set of expectations in the Guidelines would apply only if the FI used AI as an "integrated" part of its business processes.
- 3.2. In the consultation paper, MAS proposed that AI would be regarded as being used as an integrated part of an FI's business processes where either:
- (a) the lack of access to AI services or tools would disrupt workflows that the FI is materially dependent on for its business activities; or
 - (b) the AI is integrated with systems that the FI is materially dependent on for its business activities.

Proportionate application of AIRG

- 3.3. Respondents were broadly supportive of the proposal to apply the Guidelines in a proportionate manner. However, a few respondents disagreed with using operational dependence as the primary criterion to determine whether an FI is subject to the full set of expectations in the Guidelines, and

expressed concern that this framing may unintentionally capture FIs whose AI use is limited to low-risk productivity tools (e.g. Microsoft Copilot). Respondents also flagged practical challenges in assessing the extent to which AI has been integrated into workflows or systems, and suggested that the scoping should instead focus on risk materiality and impact, for example whether the failure or poor performance of AI tools or services would materially affect the FI, its customers, or its ability to meet regulatory obligations.

- 3.4. A few respondents disagreed with the classification of an internal IT helpdesk chatbot as an AI use case that is an integrated part of an FI's business process in the Annex of the Guidelines. Respondents also requested additional examples to guide FIs in their assessment of whether their AI use is considered as an integrated part of FI's business process to ensure consistent application of the AIRG.
- 3.5. A few respondents requested flexibility for FIs to scale the extent of oversight and AI risk management systems, policies and procedures proportionately based on the risk materiality of each individual use case, rather than applying all controls uniformly across all integrated AI use cases. They highlighted that not all AI use cases should attract the same level of controls, and more extensive controls should be reserved for AI use cases assessed to be of high risk materiality. Respondents also suggested the degree of reliance on AI should be a factor in determining the extent of controls to be applied for an AI use case.

MAS' Response

- 3.6. MAS agrees that the workflow-based criterion to determine whether an FI is subject to the full set of expectations in the AIRG may unintentionally capture FIs whose AI use is limited to only assistive AI tools where poor performance or unavailability of such tools is unlikely to have a material adverse impact on the FI, its customers or other stakeholders.
- 3.7. Under the revised approach, an FI may apply basic AI governance policies and procedures to oversee its AI use, if the 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. Otherwise, the FI should implement the full set of expectations set out in Sections 3 to 6 of the Guidelines. These changes are reflected within Section 2 of the Guidelines.
- 3.8. MAS has also incorporated examples of assistive AI use cases where an FI may implement basic AI governance policies and procedures within paragraph 2.4. of the Guidelines. These examples include using AI to assist in drafting or proofreading emails, summarising documents or meeting notes, or generating charts and visualisations for internal reference. Humans involved in these AI use cases would



still be reviewing and checking the outputs before using them. FIs may engage MAS bilaterally where clarification is needed on the application of Section 2 of the Guidelines.

- 3.9. MAS further clarifies that Sections 3 to 6 of the Guidelines should be implemented in a risk proportionate manner. As set out in paragraph 2.2. of the Guidelines, an FI should establish a framework for oversight, including identification, inventorisation and risk materiality assessment of AI use cases. The extent of oversight and the design of AI risk management systems, policies and procedures may be proportionate to the size and nature of the FI's activities, its use of AI, and risk profile of its AI use cases.
- 3.10. At the use case level, an FI should conduct a structured risk assessment across multiple dimensions, including the complexity of the AI technology and the degree of reliance on AI for the decision or output. Based on this assessment, the FI has the flexibility to calibrate the extent of lifecycle controls to be applied. For example, controls relating to transparency, explainability, and fairness should be applied with greater rigour for use cases of higher risk materiality, such as credit scoring, insurance underwriting, and the provision of financial advisory or fund management services, where failure, malfunction, or poor performance of AI could have potentially significant consequences on the FI's customers. Further details on the risk materiality assessment are set out in paragraph 4.10. to 4.13. of the Guidelines.

Scope of basic AI governance policies and procedures

- 3.11. A few respondents asked whether the expectation to institute basic AI governance, policies and procedures applies to an FI that has not adopted AI and highlighted such an expectation would not be commensurate with the level of AI adoption of such FIs.
- 3.12. Two respondents suggested replacing the term "guidelines" with the term "policies" when referring to an FI's basic AI governance policies and procedures in the Guidelines, to avoid confusion with the Guidelines themselves.
- 3.13. Respondents sought greater clarity on the proposed scope of basic policies originally set out in paragraph 5 of the Annex to the Guidelines, including:
- (a) Senior management designation for AI oversight: whether an FI may leverage on existing governance structures for AI oversight instead of relying on a specific senior management member;



- (b) Coverage of internal AI policies (e.g., allowed and disallowed uses of AI, requirement for human review and validation of all AI outputs before use, maintaining a list of approved AI tools, process for requesting approval of new AI tools): whether an FI has flexibility to determine what constitutes allowed and disallowed uses of AI, and the circumstances under which human review and validation of AI outputs are required;
- (c) Frequency of review: whether an FI has flexibility to determine an appropriate frequency for review of its guidelines, based on factors such as changes in AI use cases, risk profile, or business model, instead of a fixed annual review.

MAS' Response

- 3.14. As stated in paragraph 1.5. of the Guidelines, the expectations set out in the Guidelines are relevant to all FIs deploying AI. For an FI that has yet to adopt AI, MAS expects the establishment of basic AI-related policies, including policies that clearly set out whether and how AI may be used by staff, to manage risks arising from any unauthorised use of AI (such as “shadow AI”), and to facilitate a proportionate scaling-up of governance in the event of future adoption.
- 3.15. MAS agrees with the suggestion to replace the term “guidelines” used to refer to an FI’s basic AI governance policies and procedures with the term “policies” for greater clarity, and has amended paragraph 2.5. of the Guidelines accordingly.
- 3.16. MAS’ responses to the respective feedback on the scope of basic AI governance policies and procedures are as follows:
- (a) Senior management designation for AI oversight: MAS recognises that FIs may adopt different approaches for AI oversight. An FI may leverage existing governance structures for AI oversight, provided there is clear accountability for AI oversight within the organisation.
 - (b) Coverage of internal AI policies: The examples of permitted and prohibited uses of AI cited in the revised paragraph 2.5.(b) are illustrative. An FI may tailor the areas covered in its internal policies based on its own governance arrangements, the nature and risks of its AI use cases and its risk appetite. An FI has flexibility to determine the circumstances under which human review and validation of AI outputs are required.
 - (c) Frequency of review: MAS has amended paragraph 2.5.(f) of the Guidelines to provide an FI with the flexibility to determine an appropriate frequency for the periodic review of its basic AI governance policies and procedures. Reviews may also be conducted on a trigger basis such as when there are material changes to the FI’s AI use. Nonetheless, given the rapid pace of development and adoption of AI technology, MAS encourages FIs to conduct at least an annual

review of their basic AI governance policies and procedures as good practice. Where no updates are required following a review, an FI should document this accordingly.

4. Responsibilities of the Board and Senior Management

- 4.1. MAS sought comments on the proposed responsibilities of the FI's board and senior management in overseeing AI risk management. Such responsibilities included maintaining effective oversight of risks from AI use, fostering the appropriate risk culture for the use of AI, ensuring that risk management frameworks adequately identify, assess and address risks posed by AI, and ensuring that the FI's use of AI would not conflict with its ability to meet other supervisory expectations.

Discharge of responsibilities by other committees or structures

- 4.2. Respondents were broadly supportive of the proposed responsibilities of the FI's board and senior management. A few respondents requested guidance on the extent to which an FI which is part of a global group could rely on global or regional-level governance structures to fulfil the proposed responsibilities, to avoid duplication of governance structures. A few respondents sought clarity whether such responsibilities could be discharged via existing Board committees or delegated to existing individuals or committees which had the relevant expertise for AI risk management.

MAS' Response

- 4.3. An FI that is part of a global group may discharge board responsibilities either through a board-level committee delegated by the board of directors, or through a governing body or committee beyond local management (such as a regional or global committee) that exercises oversight of the risks from AI use relevant to the FI in Singapore. For other FIs, board responsibilities may be discharged via a board-level committee. A board may delegate certain responsibilities to senior management committees, but it retains overall accountability.
- 4.4. As stated in paragraph 1.5. of the Guidelines, an FI which is part of a global group may leverage the AI risk management frameworks of their parent entity to meet the expectations in the Guidelines.

Notwithstanding this, local senior management should still have adequate visibility into how the risks from AI use relevant to Singapore operations are being managed and remains accountable for ensuring that group frameworks adequately address local considerations and regulatory requirements. An FI should be able to demonstrate how such oversight is discharged, including where reliance is placed on group-level frameworks and structures.

Responsibilities of board versus senior management

- 4.5. A few respondents requested greater delineation of the expected responsibilities of the board versus that of senior management. It was highlighted that certain aspects of the Guidelines referred to the board and senior management collectively without distinction.
- 4.6. A few respondents also sought clarification on the expectations for the board (or a committee delegated by it) to approve key frameworks, structures, policies and procedures. In particular, one respondent suggested allowing the board to delegate approval of technical procedures to senior management or a specialised committee.

MAS' Response

- 4.7. MAS has further streamlined the Guidelines to enhance the clarity of the respective responsibilities of the board versus that of senior management. The board will be responsible for areas such as approving the overall governance approach and key frameworks, and ensuring that risks from AI use are addressed in the FI's risk appetite framework. Senior management will be responsible for effective implementation of the AI risk management framework, instituting appropriate controls to manage risks from AI use, and establishing clear roles and responsibilities for AI risk management across the FI. Collectively, the FI's board and senior management should maintain effective oversight of risks from AI use.
- 4.8. MAS only expects the board (or a committee delegated by it) to approve key frameworks, structures, policies and procedures which set out the overall governance approach to AI risk management. Such documents could cover areas such as the definition and scope of AI, governance and oversight structures for AI risk management, and setting the overarching tone and guiding principles for use of AI across the FI. Operational or technical frameworks, structures, policies and procedures, which may cover areas such as the FI's approach to risk materiality assessment, or the application of lifecycle controls, may be approved by senior management or relevant committees as appropriate.



Further clarifications on board responsibilities

4.9. A few respondents sought greater clarity on specific board responsibilities, including:

- (a) Technical competence: The extent to which the board should possess technical competence, receive specialised training, or rely on experts within or outside the FI, to ensure adequate understanding of AI.
- (b) Risk appetite: Guidance on incorporating risks from AI use within the FI's risk appetite framework, including clarification on what would be considered acceptable risk appetite statements or measures, and the specific expectations on the board given that this technical process of defining specific risk appetite statements or measures was usually handled by the FI's functional risk management forums.
- (c) Other responsibilities: Whether certain other responsibilities not mentioned in the Guidelines should be explicitly assigned to the board, such as setting the strategic direction of AI use or approving use cases with high risk materiality.

MAS' Response

4.10. MAS' responses to the respective items are as follows:

- (a) Technical competence: The board is not expected to possess hands-on technical expertise to challenge AI use at a granular level. However, it should maintain a sufficient understanding of risks arising from the FI's adoption of AI to provide effective oversight. This is consistent with how boards oversee other complex or technical risks, such as technology risk.

The board should be able to provide effective challenge on matters such as senior management's assessments of material risks from AI use, the adequacy of the overall control environment, and whether the FI's use of AI is consistent with the FI's risk appetite and strategic objectives. To meet this expectation, the board may determine the most appropriate approach, which may include having relevant expertise on the board itself and/or supplementing its collective expertise through other means, such as regular training or independent expert input.

While the board may receive advice from experts within or outside the FI, such advice is not a substitute for the board's own oversight responsibilities. The board remains accountable for exercising its own judgement in discharging its oversight role.

- (b) Risk appetite: MAS' intent is for the risk appetite framework to address AI-specific risks. An FI has flexibility to determine how best to incorporate such risks into its existing risk appetite framework.

The Guidelines have been updated to provide illustrative examples of qualitative statements or quantitative measures that an FI may consider. An FI's functional risk management forums may propose approaches for incorporating risks from AI use into the FI's risk appetite framework, but it is for the board to determine what is most appropriate.

- (c) Other responsibilities: MAS agrees that the board should set the overarching strategic direction for AI use, as this will drive the nature and extent of risks to which the FI will be exposed from its use of AI, and has updated the Guidelines accordingly. On approving high materiality use cases, an FI has flexibility to determine whether such approvals should rest with its board or senior management, given the operational nature of such decisions.

Three lines of defence and role of Audit

- 4.11. A few respondents sought guidance on the application of the 'three lines of defence' structure to the assignment of roles and responsibilities in AI risk management, including examples of roles and responsibilities envisaged from each line of defence. Two respondents suggested more explicitly setting out expectations on assurance provided by internal or external auditors, to identify gaps in an FI's AI risk management framework.

MAS' Response

- 4.12. The 'three lines of defence' structure can be a beneficial reference when establishing roles and responsibilities for AI risk management. FIs should adopt governance structures suited to their business context, provided these ensure clear accountabilities and appropriate segregation of duties. The Guidelines have been amended under footnote 19 to give some illustrative examples of the potential roles and functions of each line of defence, including the role of internal audit as the third line of defence under a 'three lines of defence' structure.

5. Establishing a dedicated cross-functional committee

- 5.1. MAS proposed that where the overall AI risk exposure of the FI is material, the FI should establish a dedicated cross-functional committee to ensure adequate AI oversight and proactively address potential gaps in risk management coverage.

Need for a dedicated cross-functional committee

- 5.2. While respondents broadly agreed on the need for robust cross-functional oversight of AI, views were mixed on whether a cross-functional committee was necessary to ensure such oversight. In particular, some respondents expressed concerns with the expectation for the committee to be dedicated (i.e. with sole focus on AI oversight), as this may be duplicative with existing governance structures, some of which may already involve cross-functional committees. It was suggested to give the FI flexibility to determine appropriate structures for ensuring adequate AI oversight, whether via existing cross-functional committees, or expanding the terms of reference of existing specialised committees (e.g. technology risk, model risk) to cover risks from AI use.

MAS' Response

- 5.3. MAS has amended paragraph 3.3. of the Guidelines to recognise that FIs have the flexibility to adopt varying approaches to AI risk management. Such approaches include creating new centralised functions dedicated to risks from AI use, or managing the incremental risks from AI use via existing risk management functions. Regardless of the operating model, the FI should still ensure that there is a consolidated view of risks from AI use across the organisation and that such risks are managed in a consistent and coordinated manner across all relevant functions, whether via a dedicated cross-functional committee, or other appropriate structures. An FI should be able to demonstrate how its chosen approach achieves this objective.

Expectations on the cross-functional committee

- 5.4. Some respondents proposed that the Guidelines recognise that the dedicated cross-functional committee could be established at global or regional level, to ensure consistency in AI risk management across the entire organisation, and to avoid creating duplicative structures at the Singapore level.
- 5.5. A few respondents requested greater clarity, as well as provided suggestions, on the expected responsibilities of the dedicated cross-functional committee. Suggestions fell largely into two groups:
- (a) Strategic responsibilities, such as oversight, ethical alignment, risk appetite, and approval of policies and procedures.
 - (b) Operational responsibilities, such as approving certain AI use cases, monitoring key risk indicators, reviewing AI incidents, and providing relevant reporting to the Board and senior management.



MAS' Response

- 5.6. MAS agrees that for an FI which is part of a global group, such dedicated cross-functional committees could be set up at the global or regional level. However, the FI should be able to demonstrate how the dedicated cross-functional committee ensures consistent and coordinated management of risks from AI use relevant to Singapore operations.
- 5.7. Regarding the specific responsibilities of the dedicated cross-functional committee, an FI should refer to paragraphs 3.4. and 3.5. of the Guidelines on the responsibilities of the FI's board versus that of senior management. For example, certain responsibilities such as risk appetite and approval of key policies and procedures should rest with the FI's board. An FI has the flexibility to determine which of the responsibilities stated in paragraph 3.5. of the Guidelines should be discharged by the dedicated cross-functional committee, and which should be discharged by other functional committees or governance bodies within the FI.

6. AI Identification & Inventory

- 6.1. MAS had proposed for an FI to establish systems, policies and procedures to ensure the consistent identification of AI use across all relevant business and functional areas. Clear definitions, criteria and processes, supported by robust systems, should be implemented to facilitate this identification process.
- 6.2. MAS also proposed for an FI to establish and maintain an accurate and up-to-date inventory of AI use cases, systems or models across the FI. While an FI could enhance existing inventories to include AI use cases, systems or models or establish a dedicated inventory for AI use cases, the FI should establish clear linkages between the AI inventory and other relevant inventories in the FI.

Scope of AI to be identified & incorporated into inventory

- 6.3. While respondents generally agreed with the principle of comprehensively identifying AI use, a few respondents expressed reservations over an FI's ability to practically implement processes that achieve coverage of all AI use. In particular, challenges with identifying two types of AI were highlighted:
- (a) Embedded AI: AI incorporated within, or used in the delivery of, third-party services to the FI, where the services are not explicitly marketed or sold as AI. Challenges to AI identification could arise where providers incorporate or use AI within third-party services without disclosure or prior notification to the FI.



(b) Shadow AI: Use of public AI tools by personnel in the FI. Challenges to AI identification could arise from the lack of formal approval or visibility over such tools at the organisational level.

6.4. Similarly, some respondents expressed concerns over the expectation to incorporate all AI use cases into the AI inventory. It was highlighted that such an expectation could capture low-risk or assistive AI tools which do not materially impact the FI, leading to unnecessary operational burden. It was also highlighted that a comprehensive inventory may not be achievable once the use of AI became ubiquitous in the FI. A few respondents also sought clarity on whether AI in the pilot or proof-of-concept stages should be captured in the AI inventory.

MAS' Response

6.5. The intent of the AI identification process is to ensure visibility over how AI is used within the FI, as part of ensuring effective governance, risk management standards, and controls. The challenges associated with identifying embedded AI or shadow AI do not by themselves justify excluding embedded AI or shadow AI from the AI identification process in all instances. The assessment of the risks posed by embedded AI or shadow AI can only be made by first identifying the sources and nature of such embedded AI or shadow AI. Accordingly, MAS' intent is for the AI identification process to consistently identify AI use across all relevant business and functional areas.

6.6. While an FI should endeavour to identify all instances of embedded AI, MAS recognises that the FI may face challenges in doing so. The Guidelines have therefore been revised to clarify that the identification process should minimally identify embedded AI in services provided by material third-party service providers. This reflects the need for an FI to maintain visibility over material AI dependencies that may not be immediately apparent, enabling the FI to identify and manage the risks arising from such dependencies.

6.7. MAS also recognises that it may not be practically possible to ensure complete identification of all AI use in practice. The Guidelines have been amended to state that where such constraints exist, the FI should assess the risks arising from gaps in visibility and put in place mitigating measures. These could include establishing clear policies on permitted and prohibited uses of AI; implementing technical controls to restrict access to, or the sharing of sensitive information with, public AI tools where appropriate; monitoring for unauthorised AI use; and periodically reinforcing staff awareness of these requirements.

6.8. All AI use cases should be recorded in the AI inventory to the extent possible or practicable. The intent of the AI inventory is to serve as the basis upon which an FI performs its risk assessment and applies the appropriate risk management controls across the AI life cycle. An FI can implement the AI inventory

in a risk-proportionate manner to manage the operational burden associated with implementing and maintaining the AI inventory.

Implementation of “robust systems” for identifying AI

- 6.9. A few respondents requested clarity on the expectation that “robust systems” be implemented to support the process of AI identification. In particular, how such expectations should be guided by the principle of proportionality, and whether such systems should include tools which conduct automated surveillance to discover shadow AI or un-authorised AI use. A few respondents expressed concerns that such surveillance tools may be invasive in nature, and disproportionate to the risks posed.

MAS’ Response

- 6.10. MAS’ intent is that such systems facilitate the process of AI identification, which may otherwise be an uncoordinated, laborious and highly manual effort for an FI which has a significant number of AI use cases or more mature adoption of AI. The Guidelines have been amended to provide an FI flexibility in determining how appropriate tools and systems may be implemented to facilitate the identification process. For example, an FI with a lower number of AI use cases may consider it sufficient to rely on manual processes for staff to identify and attest to their use of AI. In contrast, an FI with a higher number of AI use cases may deem it necessary to employ more sophisticated systems to ensure that AI is consistently identified across the organisation. Such systems may also facilitate the capture of AI use cases, including relevant attributes, into the AI inventory.
- 6.11. An FI has the flexibility to determine whether the application of automated surveillance or other discovery tools to detect unauthorised AI use is appropriate in achieving the overarching intent of consistent identification of AI. The feasibility of such tools may also be evaluated against other measures that the FI may put in place to mitigate risks associated with un-identified use of AI across the FI (as stated in paragraph 4.4. of the Guidelines).

Implementation of the AI inventory

- 6.12. Some respondents requested for more clarity on the attributes which should be captured in the AI inventory, including whether MAS had any expectations on attributes which must minimally be captured. Respondents also suggested various attributes for MAS’ consideration, such as data sensitivity classifications or names of third-party vendors providing the AI model or service.

- 6.13. A few respondents requested for clarity on the level of granularity expected of the AI inventory. This included specifying whether inventory entries should be at the level of the AI use case or process, or at the model or product level. A few respondents suggested that MAS allow for aggregate or 'category-level' registration for low-risk AI, such as a single entry for a specific product (e.g., Microsoft Copilot) or category (e.g., General Productivity Tools), to mitigate the operational burden of exhaustively documenting AI use cases across the relevant tool or category. A few respondents also requested flexibility to adopt tiered inventories with differing levels of detail across AI, such as by capturing detailed attributes only where the AI use case was material.
- 6.14. One respondent proposed for flexibility in the underlying tool or system used to support the AI inventory (i.e. whether in a dedicated system or via manually updated spreadsheets).

MAS' Response

- 6.15. The intent of the AI inventory is to allow an FI to maintain sufficient visibility of its AI adoption and facilitate effective risk management. The level of granularity and details to be collected will vary depending on the nature of the AI use case, the level of AI adoption across the FI, and the risk materiality of the use case. For example, an FI which deploys complex, multi-agent AI systems in material use cases may reflect attributes such as agent-specific identifiers and information on the tools and systems the agent can access, and to document such dependencies at the agent level. Correspondingly, an FI which primarily uses AI for assistive purposes on low-risk tasks may not find such attributes necessary.
- 6.16. In response to the feedback, the Guidelines have been amended in paragraph 4.7. to clarify expectations on establishing an appropriate granularity for the AI inventory. An FI may also determine whether it is necessary to invest in tools and systems that facilitate the implementation of the AI inventory, such as features for automated tracking of approvals and issues, identification of interdependencies between models, and lifecycle tracking with checkpoints at key risk management stages. Ultimately, the FI should be able to demonstrate that its AI inventory enables it to maintain an accurate and up-to-date view of its AI use cases.

Frequency of updating the AI inventory

- 6.17. A few respondents requested clarity on the expected frequency for updating the AI inventory. It was highlighted that maintaining a real-time inventory may be resource-intensive, especially for an FI engaging in rapid AI development. Respondents also suggested that inventory updates be part of change management, i.e. triggered by changes in the model or its status.

MAS' Response

- 6.18. An FI has the flexibility to determine a frequency for updating its AI inventory that is appropriate in its own context, taking into account the risks to which it is exposed, the nature of AI used, and the attributes captured in its AI inventory. MAS will thus not prescribe a fixed update frequency, and an AI inventory maintained in 'real-time' may not be necessary in all cases. That said, an FI should assess whether more frequent updates are warranted for AI deployed in material use cases, or for AI that undergoes dynamic or frequent changes.
- 6.19. In addition to periodic reviews, updates to the inventory may be triggered by events, such as material changes to the AI use case, or the systems and models that support the use case. Updates may also be appropriate where the risks associated with an AI use case evolves over time. For example, the impact of the AI use case may increase as the FI's operations expands, while concentration risks may change as other AI use cases relying on the same system or provider are introduced or decommissioned. An FI should therefore ensure that its AI inventory is reviewed and updated in a manner that remains fit for purpose, i.e. to allow the FI to capture changes in the use of AI to enable appropriate oversight and facilitate effective risk management.

Linkages between AI inventory and other inventories

- 6.20. A few respondents requested for more clarity on the expectation that clear linkages be established between the AI inventory and other relevant inventories in the FI. Respondents requested for examples of which attributes in the inventory should be linked, as well as which inventories the AI inventory should establish linkages to. One respondent expressed concern that such linkages could become complex due to many-to-many relationships between inventories, as well as difficult to implement given lack of transparency from third-party AI providers.

MAS' Response

- 6.21. The intent of establishing clear linkages between the AI inventory and other relevant inventories is to support effective and coherent oversight of AI-related risks, by enabling an FI to draw meaningful connections across related information that is already maintained elsewhere. MAS does not expect an FI to establish linkages with all its other inventories. Rather, an FI should exercise judgement in determining which linkages are necessary and useful, having regard to the structure of its AI inventory and how AI-related risks are monitored and managed in practice. The Guidelines have been amended to more clearly specify the purpose of such linkages.

- 6.22. For example, where an FI maintains a register of third-party arrangements alongside a central AI inventory covering use cases that rely on third-party AI providers, establishing linkages between these inventories (e.g., through consistent vendor or service identifiers) can facilitate a consolidated view of the FI's relationship with the vendor, including risks arising from the use of third-party AI and aggregate reliance for concentration risk assessment.
- 6.23. Additionally, an FI that incorporates AI use cases into existing data or IT asset inventories may apply consistent data or system identifiers across these inventories, to enable a more integrated assessment of the datasets and systems underpinning each AI use case as part of its risk materiality assessment.

7. Risk materiality assessment

- 7.1. MAS had proposed that an FI establish an assessment methodology to evaluate the risk materiality of an AI use case, system or model, which should be applied consistently to perform risk materiality assessments for each AI use case, system or model used by the FI. The assessment should take into account both the inherent risk materiality of an AI use case, system or model before the appropriate risk management controls are applied, as well as the residual risk materiality after risk management controls are applied.
- 7.2. While an FI had flexibility to consider various risk dimensions in the risk materiality assessment methodology according to its context, MAS proposed that the assessments should minimally cover three dimensions: Impact, Complexity, and Reliance.

Level of application and risk dimensions

- 7.3. A few respondents proposed that risk materiality should focus on how AI is used, rather than the system or model in isolation. Respondents highlighted examples where simple models that can materially impact clients should be assessed as higher risk than more complex models that are not customer-facing.
- 7.4. Respondents also suggested a wide range of additional risk dimensions to be integrated with, or explicitly added to, the proposed risk dimensions of Impact, Complexity and Reliance. Common suggestions included: data privacy, autonomy, level of interconnectedness or shared dependence on common models or systems, extent of third-party dependence, reversibility of AI-driven decisions, stability of AI outputs, and explainability.



MAS' Response

- 7.5. MAS agrees that the assessment of risk materiality should be anchored to the AI use case. To allow for the application of appropriate lifecycle controls, risk materiality assessments should be holistic and factor in various risk dimensions, including the potential consequences of failure, malfunction, or poor performance of the AI, and not the technical complexity of the underlying systems or models alone.
- 7.6. For example, where the same model supports two AI use cases, but is deployed in an automated manner in one, and with full human-in-the-loop oversight in another, this difference may be reflected in the respective risk materiality assessments of the two use cases. The Guidelines have been amended accordingly under paragraph 4.10., including relevant references to the application of risk materiality assessments under the section on lifecycle controls in Section 5.
- 7.7. MAS has considered the additional risk dimensions proposed by respondents, and amended the Guidelines to incorporate explainability and third-party considerations within the definitions of Complexity and Reliance. Beyond the three risk dimensions, MAS does not intend to prescribe an exhaustive list of risk dimensions for AI risk materiality assessments. FIs may adopt or retain existing frameworks or approaches that incorporate additional dimensions into the risk materiality assessment as appropriate, and are encouraged to consider additional risk dimensions beyond the proposed dimensions of Impact, Complexity and Reliance. For instance, an FI may wish to consider the extent of access to tools and systems granted as an additional dimension for AI risk materiality assessments of agentic AI, given that agents capable of taking actions across multiple tools and systems may compound data compromise risks as their permissions expand.

Assessment of Complexity

- 7.8. Some respondents expressed concerns over the inclusion of Complexity as a dimension to be considered in risk materiality assessments. They argued that Complexity does not necessarily correlate with risk, citing examples where technically sophisticated models may pose lower risks when used in low-risk use cases, or where simpler models may pose higher risks in more sensitive contexts. Other concerns related to the subjective and evolving nature of Complexity, and practical challenges in assessing Complexity early in the lifecycle or where third-party models limit transparency.



MAS' Response

- 7.9. MAS has considered the feedback and has decided to retain Complexity as a risk dimension. Complexity is not intended to be assessed in isolation, nor does it determine risk materiality on a standalone basis. Rather, Complexity is one of several dimensions that should be considered in combination, alongside factors such as Impact, when determining the overall risk materiality of an AI use case. As clarified in paragraph 7.19., FIs have flexibility in determining how individual risk dimensions are combined or weighted.
- 7.10. MAS recognises that assessments of Complexity may involve judgement and evolve as an FI's understanding of an AI technology develops. This does not diminish the relevance of Complexity as a risk dimension. Rather, it underscores the importance of assessing Complexity in the context of the FI's own understanding of, and ability to assess, the technology.
- 7.11. Accordingly, an FI that is less familiar with an AI technology, or has more limited ability to understand and assess its risks, may reasonably assess Complexity as higher and apply correspondingly more rigorous lifecycle controls, such as enhanced testing and pre-deployment review. Where the use of third-party AI limits an FI's visibility into the nature or risks of the AI technology, the FI should similarly consider whether this limitation warrants a higher assessment of Complexity. An FI may revise its assessment of Complexity, and the corresponding lifecycle controls, as its familiarity with or visibility into the AI technology develops over time.

Assessment of Reliance

- 7.12. A few respondents raised concerns or sought clarification on the relevance of considering the availability of alternatives in assessing Reliance. One respondent highlighted that the availability of alternatives may be interpreted narrowly as theoretical substitutability rather than operational feasibility under normal control and volume conditions.

MAS' Response

- 7.13. MAS has decided to remove the reference to the availability of alternatives as a sub-dimension of Reliance, and to clarify that assessments of Reliance should focus on the extent to which an AI use case depends on AI-generated decisions or outputs. Relevant considerations include the level of autonomy granted to the underlying AI system or model, the degree of human oversight, and the extent of dependence on third-party providers if relevant.

7.14. The removal of the availability of alternatives as a sub-dimension of Reliance does not mean that fallback options are irrelevant to AI risk management. As set out in paragraph 5.3. of the Guidelines, an FI should develop contingency plans for AI use cases of high risk materiality. Such contingency plans should outline fallback options that will be operationally feasible under a range of operating conditions, such as alternative systems or manual processes, to ensure business continuity in case of AI failure or unexpected behaviour.

Determination of risk materiality ratings

7.15. A few respondents requested for guidance, or suggested various approaches and examples, on how to assess the level of risk under Impact, Complexity and Reliance, with a view to promoting consistent application across FIs. For example, respondents sought clarification on what level of reliance would be considered as “high Reliance”.

7.16. Respondents also requested guidance on how assessments of Impact, Complexity and Reliance should be combined when deriving an overall risk materiality rating. Some respondents suggested approaches where certain dimensions were considered of primary importance, with other dimensions influencing the overall risk materiality only under certain conditions (e.g., considering Complexity and Reliance only where Impact is assessed to be higher).

7.17. Lastly, a few respondents requested clarity on what AI would be deemed as “high” risk materiality, noting that certain controls in the Guidelines apply where risk materiality is assessed to be high.

MAS’ Response

7.18. MAS has provided illustrative use cases and characteristics for each risk dimension in the table below to support risk materiality assessments. These illustrations are intended to be read as relative comparisons (i.e. higher vs lower Impact, Complexity or Reliance) and do not imply or prescribe any specific absolute rating. An FI should exercise judgement in determining the relevance of the illustrations to its specific circumstances.

Dimension	AI which could be assessed of higher risk under the dimension	AI which could be assessed of lower risk under the dimension
Impact	Use of AI for credit scoring Poor performance of AI is likely to have material adverse impact on	Use of AI for personal productivity (e.g. internal document summarisation)

	customers' access to and pricing of credit products. AI processes personal data which could be of a sensitive nature.	Poor performance is unlikely to have a material adverse impact on customers or critical business functions.
Complexity	Use of Generative AI systems provided by third-parties - Approaches for testing, monitoring and explaining outputs are still evolving. - FI may have limited visibility or understanding of the design or training data of the underlying model.	Use of regression-based AI models developed internally by the FI - Methodology is well-understood and explanation of outputs is straightforward. - FI has full visibility into the design, training approach and data.
Reliance	Use of AI for automated trade execution - Decisions are executed in an automated way with limited human review of individual transactions. - Decisions are difficult to reverse once executed.	Use of AI to assist in recommending responses to consumer queries - Staff review each recommendation before responding. - Poor recommendations can be rectified by staff, reducing impact on consumer outcomes.

7.19. MAS does not intend to prescribe a uniform or formulaic approach for combining individual assessments of Impact, Complexity and Reliance into an overall risk materiality rating, nor to define a universal threshold for “high” risk materiality. Such assessments depend on an FI’s specific AI use cases, risk profile and operating context. Accordingly, an FI has flexibility to determine how the respective risk dimensions are weighted or combined in its overall risk materiality assessment. Such flexibility extends to approaches where certain dimensions are considered of primary importance, while others only influence the overall risk materiality ratings under certain conditions.

Inherent versus residual risk materiality assessments

7.20. One respondent sought clarity on approaches for determining residual risk when assessing AI risk materiality. Another respondent asked whether the requirements that apply to AI of high risk materiality (e.g., on independent validation and developing contingency plans) should be triggered based on inherent or residual risk.



MAS' Response

7.21. The requirements specific to AI use cases of high risk materiality, including more rigorous lifecycle controls, are to be applied based on inherent risk materiality. The intent is to ensure that AI use cases with high inherent risk are subject to appropriate controls upfront and on an ongoing basis, such that residual risks are adequately mitigated and remain within the FI's risk appetite.

7.22. MAS will not prescribe a specific methodology for assessing residual risk materiality. An FI has flexibility to apply an approach consistent with its existing frameworks and methodologies, including assessing residual risk at the level of individual risk dimensions before aggregating them at the use case level, or adjusting inherent risk materiality at the use case level to reflect the effect of controls.

8. Roles and responsibilities for AI risk management systems, policies and procedures

8.1. MAS proposed that an FI should assign clear roles and responsibilities for AI identification, inventorisation of AI, and AI risk materiality assessment. MAS also proposed that an FI should designate a control function to be responsible for each process:

- (a) AI identification: the designated control function should be responsible for areas such as ensuring the consistent application of the identification process across the FI, setting up attestation processes, or acting as the final arbiter in determining whether AI is being used.
- (b) AI inventory: the designated control function should be responsible for areas such as policies and procedures relating to the inventory, maintenance and update of the inventory, attestation process, and regular reviews of the scope of the inventory.
- (c) Risk materiality assessment: the designated control function should be responsible for areas such as setting up attestation processes, ensuring clear documentation is maintained, and acting as the final arbiter in determining risk materiality.

Designation of a control function

- 8.2. A few respondents requested flexibility to designate existing control functions, including functions established at the Group-level which had Group-level oversight responsibilities. It was highlighted that requiring new or standalone control functions at the Singapore-level could be duplicative in nature.
- 8.3. A few respondents also requested flexibility to adopt federated approaches, where frameworks and policies are facilitated by the designated control function, and specific components are owned and implemented by different relevant stakeholders.

MAS' Response

- 8.4. An FI may opt to designate an existing function, including those with Group-level oversight responsibilities, so long as the FI is able to demonstrate adequate oversight on AI used in Singapore. For example, the FI should be able to demonstrate that the designated control function ensures consistent application of the identification process with respect to AI used in Singapore, and appropriately captures attributes of such AI in its AI inventory.
- 8.5. MAS recognises that AI identification, inventorisation and risk materiality assessments may be carried out by relevant business or functional units outside of the designated control function. The Guidelines have been amended to make it clearer that relevant business or functional units may execute these processes, so long as the designated control functions retain sufficient and independent oversight in the execution of such processes, including to facilitate consistent application of the processes across the FI.

Setting up attestation processes

- 8.6. A few respondents requested for clarity on setting up attestation processes, such as whether such attestations related to the assigned roles and responsibilities or the application of AI identification.

MAS' Response

- 8.7. The intent of such attestation processes is to provide assurance that:
- (a) AI use within the FI is appropriately identified in a timely manner;



- (b) the FI's AI inventory remains accurate and complete, reflecting new or updated AI use cases; and
- (c) risks arising from the FI's AI use cases are appropriately assessed and kept under review, in accordance with the FI's policies and procedures.

8.8. MAS recognises that FIs may have varying approaches to achieve these outcomes. The Guidelines have thus been amended to focus on the outcomes to be achieved, rather than to require the establishment of specific attestation processes. An FI has flexibility in approach (including whether to use attestations or other mechanisms) but should be able to demonstrate to MAS how its chosen approach achieves these outcomes.

9. Third-party AI management

9.1. MAS had proposed that an FI should ensure that onboarding, development and deployment controls for third-party AI are adequate, having regard to the risks posed by the use of third-party AI. The FI should also ensure that AI developed by third parties are subject to appropriate reviews, as well as establish processes to receive notifications of updates or changes to third-party AI and manage and assess the impact of such updates or changes. The FI should consider various areas, such as the level of transparency from third-party AI providers and updating legal agreements to facilitate clear expectations and responsibilities on the third-party AI provider.

Division of responsibilities between the FI and third-party providers

9.2. A few respondents sought greater clarity on the extent to which an FI should be held responsible for aspects of AI risk management that were controlled by the third-party AI provider. It was suggested that the third-party AI provider, rather than the FI, remain responsible for areas such as model safety and security, training data quality, model architecture, infrastructure resilience, platform security, vulnerability management, and compliance with applicable technical standards.

MAS' Response

9.3. The choice of whether to onboard or use a third-party AI service rests with the FI. As such, even where an FI may not be in a position to influence or compel third-party AI providers to employ robust AI risk

management practices, the FI remains accountable for how third-party AI is used within its organisation, including the possible impact on the FI, its customers and other stakeholders.

- 9.4. Accordingly, the onus is on the FI to ensure that risks arising from its use of third-party AI are appropriately managed. In deciding whether to onboard or continue using a third-party AI service, MAS expects an FI to maintain a risk-proportionate degree of visibility into the risk management practices of the third-party AI provider to inform its decision-making.
- 9.5. The Guidelines have been amended to more clearly state that the various aspects of third-party AI risk management are to be considered and may be met to varying degrees depending on the risk materiality of the AI use case. Where an FI is unable to apply sufficient controls on third-party AI to keep the residual risk materiality of the AI use case within the FI's risk appetite, the FI should consider limiting or suspending the third-party AI service, or replacing the relevant third-party provider.

Addressing limited transparency from third-party providers

- 9.6. Some respondents highlighted that third-party AI providers typically treat various aspects of the model development process as proprietary in nature, such as details on training data and how the model has been trained and tested. Respondents added that such vendors would not likely provide sufficient transparency for the FI to assess the robustness of the model development process. Respondents requested for greater guidance on what constituted sufficient compensatory testing where transparency was limited, and how such expectations should be calibrated by risk materiality.
- 9.7. To address the issue of limited transparency, some respondents proposed that MAS consider allowing an FI to leverage standardised and independent certifications, such as ISO 42001 or SOC Type 2, as valid alternatives. A few respondents suggested that MAS consider a "safe harbour" with relaxed requirements for AI from reputable providers.

MAS' Response

- 9.8. As noted in the previous response, where an FI chooses to onboard or use a third-party AI service, the FI remains responsible for managing the risks arising from such use. MAS recognises that it may not always be possible to obtain sufficient transparency from the third-party AI provider to independently verify how key risks have been addressed in respect of the AI service provided.

- 9.9. Where such transparency is insufficient, MAS agrees that an FI may consider relying on certifications or external assessments of the third-party AI provider to inform its risk assessment and decision-making. The FI should satisfy itself that the certifications or external assessments are conducted by parties which possess the requisite knowledge and expertise, are independent of the third-party AI provider and appropriately cover the key risks for which sufficient transparency is unavailable. In contrast, relying on self-attestations or “statements” by the third-party AI provider that the risks have been addressed, particularly where unsupported by evidence, would generally not be considered as effective approaches to address limitations in transparency.
- 9.10. Certifications and external assessments of the third-party AI provider are not the only means by which an FI may address limitations in transparency; an FI may also supplement such certifications and external assessments and/or adopt other risk-proportionate measures to obtain sufficient assurance over a third-party AI provider. For example, by conducting additional testing to better understand the behaviour and limitations of the third-party AI when applied to the FI’s own data or use cases, or applying greater human oversight over the outputs of such AI. The Guidelines have been updated to set out MAS’ expectations on the use of certifications or external assessments.

Processes for updates and changes to third-party AI

- 9.11. A few respondents sought clarity on whether an FI was expected to establish processes to receive notifications of all updates or changes to third-party AI, and to correspondingly manage and assess the impact of all such updates or changes. It was highlighted that an FI may face difficulties in ensuring that the third-party AI provider appropriately provided the FI with notifications of all updates or changes made. Further, requiring an FI to assess the impact of all such updates or changes may not be risk-proportionate if the third-party AI was only being used for use cases of lower risk materiality.

MAS’ Response

- 9.12. As noted in the previous responses, where an FI chooses to onboard or use a third-party AI service, the FI remains responsible for managing the risks arising from such use, taking into account the risk materiality of the relevant AI use case.
- 9.13. MAS recognises that an FI may not always be in a position to influence or compel third-party AI providers to notify it of all updates or changes to third-party AI, and that the impact of such updates or changes will likely vary with the risk materiality of the use case. MAS thus does not expect an FI to establish processes to receive and assess all updates or changes to third-party AI in all circumstances.

- 9.14. Instead, an FI should ensure that contractual arrangements with third-party providers allow the FI to have a risk-proportionate degree of visibility into updates or changes to third-party AI that could reasonably be expected to affect the continued performance, behaviour or risk profile of the AI use case. Where relevant, such updates or changes may need to be tested to maintain satisfactory performance, and the extent of review prior to continued use should be commensurate with the risk materiality of the AI use case, and comparable to how significant changes to internally developed AI would be assessed. This is elaborated upon in paragraph 10.74. of this response paper.
- 9.15. The Guidelines have been clarified to state that the extent to which an FI is able to receive notifications of updates or changes to third-party AI post-deployment is a consideration when deciding whether to onboard or use third-party AI. This should be read alongside the lifecycle control on change management, under which paragraph 5.25.(a) has been updated to clarify that where significant changes to AI may occur without an FI's prior review, compensating controls such as enhanced monitoring should be applied to mitigate the associated risks.

Interaction with MAS' outsourcing framework

- 9.16. A few respondents sought clarification on how the Guidelines interact with existing MAS requirements on outsourcing, such as the Outsourcing Guidelines. Two respondents suggested more explicit cross-referencing to indicate that the Guidelines are intended to be read in conjunction with MAS' requirements on outsourcing.
- 9.17. A few respondents also sought clarity on the relationship between the risk materiality assessment in the Guidelines, and the assessment of the materiality of an outsourcing arrangement under MAS' outsourcing framework. One respondent suggested that the expectations in the Guidelines could create a parallel and duplicative framework, adding complexity to third-party frameworks and assessments. Another respondent asked whether providers of third-party AI would automatically be classified as providing material outsourcing arrangements under MAS' outsourcing framework.

MAS' Response

- 9.18. The Guidelines are intended to be read in conjunction with existing legislation and guidelines issued by the MAS, including on outsourcing. The Guidelines have been amended in paragraph 1.2. to make this clear, without providing an exhaustive list of all such requirements.

- 9.19. Under MAS' Guidelines on Outsourcing³, an FI assesses the materiality of an outsourcing arrangement based on the risks posed by the service provider relationship, including factors such as the importance of the outsourced activity to the FI's operations, the potential financial and reputational impact should the service provider fail to perform or suffer a security breach, and the FI's ability to maintain internal controls and meet regulatory requirements in such scenarios.
- 9.20. In contrast, the risk materiality assessment under the Guidelines is performed at the level of the AI use case. The assessment considers not only the impact of the AI use case, but also dimensions such as Complexity and Reliance. Risk materiality assessments should be applied consistently across all AI use cases, including those with no related outsourcing arrangement, such as where AI systems and models are developed internally.
- 9.21. MAS does not consider risk materiality assessments of AI use cases to be duplicative of the materiality assessments of outsourcing arrangements. For the same reason, providers of third-party AI would not automatically be classified as providing a material outsourcing arrangement to the FI.

10. Other lifecycle controls

- 10.1. MAS had proposed that an FI should plan for and implement robust controls covering the entire AI life cycle, based on relevance to the AI model, system or use case and proportionate to the assessed risk materiality. Key areas that an FI should assess for relevance and apply in a proportionate manner include data management, fairness, transparency and explainability, human oversight, selection of AI, evaluation and testing, technology and cybersecurity, reproducibility and auditability, pre-deployment reviews, post-deployment monitoring, and change management.

Data Management

- 10.2. A few respondents sought greater clarity on how the principle of proportionality would apply to data management, particularly in the context of third-party AI. They highlighted practical difficulties in assessing training data used by third-party models where third-party AI providers were unwilling or unable to provide sufficient transparency.

³ MAS has published a Consultation Paper on Proposed Guidelines on Third-Party Risk Management (TPRMG) on 6 March 2026, which is intended to supersede the Guidelines on Outsourcing. Once the TPRMG is finalised, an FI should refer to the TPRMG in reading this paragraph.

10.3. Respondents also sought guidance on specific data-related issues, including the circumstances under which customer data may be used to train AI models (e.g., whether explicit consent from data subjects is required), how expectations apply when data was transferred or shared cross-border (e.g., to train or fine-tune third-party AI models outside of Singapore), and how data representativeness should be assessed where external datasets or pre-trained models are used. One respondent suggested clarifying the definition of “full range of real-world conditions”. Another respondent asked how expectations on data management should be applied in the context of Generative AI.

MAS’ Response

10.4. As noted in earlier responses, proportionality is a key consideration under the Guidelines. An FI should calibrate the extent and rigour of data governance and management controls to the risk materiality of the AI use case. For example, use cases with higher risk materiality may warrant stricter thresholds for data quality, security and representativeness, including consideration of data across a broader range of relevant real-world conditions.

10.5. MAS recognises that limited transparency from third-party AI providers may constrain an FI’s ability to directly assess whether certain expectations in the Guidelines have been met, including those relating to representativeness, quality, and privacy of data. Paragraph 5.11. of the Guidelines has been updated to clarify that, when deciding whether to onboard or use third-party AI, an FI should consider the extent to which transparency of the third-party AI provider enables it to assess whether appropriate lifecycle controls – including on data management – can or have been applied.

10.6. Expectations on data privacy and protection, including whether consent from data subjects is required, continue to be governed by existing regulatory requirements and expectations, including those set out by the MAS and relevant regulatory agencies, such as the Personal Data Protection Commission. These requirements apply irrespective of whether data processing or model training is conducted in Singapore or cross-border.

10.7. Expectations on data representativeness apply regardless of whether external datasets or pre-trained models are used. An FI should consider the provenance of such data and assess whether it is representative of the intended context of use. Where an FI is unable to assess representativeness in respect of training data directly, it may apply compensating controls, such as additional testing of the AI on data that is representative of the FI’s use case. To allow greater flexibility, the Guidelines have been updated to clarify that representativeness should be assessed across ‘a range’ of real-world conditions, rather than ‘the full range’, in a risk-proportionate manner.

- 10.8. Expectations on data management apply similarly to Generative AI, although they may be met in different ways compared to traditional AI. For example, where an FI is unable to directly assess the quality of training data, it may rely on outcome-based testing using representative ground truths. Additional controls should also be applied to address data security risks specific to Generative AI, such as preventing inadvertent leakage of sensitive data through prompt injection.

Transparency and Explainability

- 10.9. A few respondents suggested clarifying that expectations on explainability should be context-appropriate, reflecting the position of different actors in the AI value chain. A few respondents provided suggestions, or sought guidance, on what information an FI should disclose to its customers on the use of AI. One respondent suggested that disclosures be layered in nature, with high-level disclosures provided to all customers, with details provided on request.
- 10.10. A few respondents sought clarifications on expectations or acceptable approaches for providing explainability to customers where highly complex models were used. They highlighted potential challenges in applying the same standards for transparency and explainability in respect of complex models such as large language models. In this context, two respondents proposed that outcome-based approaches, such as functional testing, should be accepted in place of feature-level explanations.
- 10.11. Two respondents suggested that explainability may not be a meaningful consideration for low-risk AI use cases, such as internal productivity tools.

MAS' Response

- 10.12. MAS agrees that transparency and explainability is highly contextual, and should be driven by the nature of the AI use case and the needs of the intended audience. An FI is therefore expected to determine what, and to whom, information should be disclosed based on these considerations and in a risk-proportionate manner.
- 10.13. Where an AI use case does not impact customers, an FI may reasonably conclude that customer-facing disclosures are not relevant. However, transparency and explainability considerations may still apply in relation to other stakeholders, such as end-users and staff in the FI. For example, technical explanations on model behaviour, limitations, and risks may be relevant for internal stakeholders involved in oversight or independent validation, especially where the AI use case is of high risk materiality.

- 10.14. Where disclosures to customers are relevant, the information provided should be aligned with their expected needs⁴. Such disclosures need not be overly technical or detailed in nature, and may instead focus on matters such as how the use of AI may affect customers, the basis on which decisions are made, and channels for further enquiries or redress. The Guidelines have been revised to provide greater clarity on these expectations.
- 10.15. MAS recognises that there may be inherent limitations in the explainability of more complex AI models. Explainability is not an end in itself but supports other outcomes, such as providing appropriate transparency on how AI systems or models produce their outputs, or enabling assessments of conceptual soundness. These outcomes may not be necessary for all AI use cases, such as internal productivity tools. Where such outcomes are necessary, an FI has flexibility in the methods or tools used to achieve an appropriate degree of explainability, taking into account the context and nature of the AI use case and its assessed risk materiality. However, where an FI is unable to achieve an appropriate degree of transparency and explainability for a given use case using a particular model, the FI should re-assess whether such a model is appropriate for that use case. This is specified under paragraph 5.12. of the Guidelines, which expects an FI to balance the objective of the AI use case against factors such as the need for transparency or explainability when deciding on the use of a given AI model.

Fairness

- 10.16. Two respondents asked how proportionality should be applied in determining the rigour or relevance of fairness assessments for a given AI use case. One respondent sought clarification on whether each AI use case is expected to explicitly define and document its own “fair” outcomes, or whether fairness principles could be established at a group or framework level and then assessed for relevance at the individual use case level.
- 10.17. A few respondents requested further guidance or examples of fairness metrics and protected attributes in the Singapore context, to be considered as part of fairness assessments. One respondent suggested to more explicitly link fairness assessments to MAS’ FEAT principles.

⁴ Relevant frameworks on transparency practices for customer-facing AI include IMDA’s Transparency Guidelines for Generative AI Chatbots, which can be accessed at <https://www.imda.gov.sg/-/media/imda/file/emerging-technologies-and-research/artificial-intelligence/transparency-guidelines-for-generative-ai-chatbots.pdf>.



MAS' Response

- 10.18. The FEAT principles are a relevant reference for defining “fair” outcomes, and includes illustrations that may be helpful in shaping fairness assessments. An FI has the flexibility to reference other frameworks or principles, provided that the FI can demonstrate how it mitigates harmful biases and discriminatory outcomes from its use of AI. The Guidelines have been amended to clarify that an FI may take reference from relevant frameworks and principles such as the FEAT principles when defining “fair” outcomes.
- 10.19. MAS considers proportionality to be primarily applicable to the rigour of fairness assessments, rather than whether fairness should be considered at all. Where fairness is a relevant consideration for an AI use case, an FI should conduct a fairness assessment, with the depth and form of that assessment calibrated to the risk materiality of the use case.
- 10.20. For example, use cases of lower risk materiality may warrant only simplified qualitative checks to verify that model outputs do not exhibit obvious biases. Use cases of higher risk materiality may require more structured assessments, such as examining the use of protected attributes or close proxies, and setting quantitative metrics and thresholds to monitor for potential discriminatory outcomes post-deployment.
- 10.21. There is no single universally applicable metric or approach for measuring fairness. Whether a specific attribute should be considered a relevant protected attribute depends largely on the FI’s use case and operating context. Accordingly, an FI may establish fairness principles or methodologies at a group or framework level, but should assess their relevance and application at the level of individual AI use cases. Such methodologies may include thresholds to identify cases where AI may be generating systematic and unjustified disadvantages to certain groups.

Selection of AI

- 10.22. A few respondents requested greater clarity on how proportionality should be applied in justifying and documenting the process of AI selection. Two respondents highlighted that requiring the selection process to be justified and documented in all cases could be overly burdensome. One respondent highlighted that it may not be practical to expect all selections to be supported by theory, research, or accepted industry practice.
- 10.23. One respondent highlighted that while the Guidelines expected developers to carry out the selection process, this may not always reflect how responsibilities are allocated in practice.



MAS' Response

- 10.24. MAS agrees that the selection process should be carried out in a risk-proportionate manner. The Guidelines have been amended to clarify that the extent to which the selection process is justified and documented should be commensurate with the risk materiality of the associated AI use case. The Guidelines have also been revised to indicate that the selection process may draw reference from theory, research or accepted industry practice, rather than requiring this in all cases.
- 10.25. Finally, the Guidelines have been amended to provide an FI with flexibility in determining the appropriate party to carry out the selection process, in a manner consistent with its internal governance and development arrangements.

Evaluation and Testing

- 10.26. A few respondents requested for further guidance on acceptable testing approaches and performance thresholds, particularly in respect of Generative AI models that may require different evaluation methods. One respondent asked whether testing at the model level alone would suffice, and whether separate use case-level testing is required where multiple use cases share the same underlying model but differ by prompts, or where prompts may be dynamic in nature.
- 10.27. A few respondents expressed concerns in conducting appropriate evaluation and testing in respect of third-party systems or models, particularly where third-party providers treat their models as proprietary and do not disclose details of the model design or training.

MAS' Response

- 10.28. Given the wide range of possible AI use cases applied across the industry, MAS does not intend to prescribe evaluation measures or testing approaches for specific types of technologies or use cases, including for Generative AI. An FI should establish evaluation measures and testing approaches that are aligned with each use case's intended objective, taking reference from relevant best practices as appropriate.
- 10.29. Similarly, MAS does not intend to set out indicative performance thresholds for each evaluation measure, as these may differ based on the risk materiality of the use case and the FI's risk appetite. For example, stricter thresholds may apply where the use case is customer-facing in nature due to the

higher impact of poor performance, or where the model will make decisions in an automated manner with limited human oversight.

- 10.30. Evaluation and testing should generally focus on the performance at the use case level rather than at the system or model level. This recognises that the same system or model may behave differently across use cases depending on how it is trained, configured or prompted. Accordingly, Generative AI models and AI agents should be assessed against prompts or tasks that are representative of their intended use cases. While evaluation need not cover every single possible prompt (e.g., where dynamic prompts are used), an FI should still conduct testing around the most likely or impactful scenarios, in a manner proportionate to the risk materiality of the use case.
- 10.31. On third-party AI, evaluation and testing should focus primarily on the performance of third-party AI in the context of the FI's use cases, including, where practicable, using the FI's own data. This is set out in paragraph 5.10. of the Guidelines. This ensures that datasets used for testing are representative of the intended use of third-party AI by the FI. That said, MAS recognises that an FI may not have such data available at the point of initial deployment. The Guidelines have thus been updated to clarify that pilots or phased rollouts may be considered as means to collect additional representative data, where relevant, to support further evaluation and testing.

Technology and Cybersecurity

- 10.32. Two respondents suggested to more clearly incorporate proportionality when setting out requirements on managing technology and cybersecurity risks, including in relation to the frequency and rigour of technology and cybersecurity reviews.
- 10.33. A few respondents sought clarity on how the technology and cybersecurity risk expectations in the Guidelines interact with other MAS regulatory requirements such as the Technology Risk Management Guidelines (TRMG), and whether existing technology risk controls established under such guidelines could be relied upon to address AI-related technology and cybersecurity risks.
- 10.34. A few respondents suggested to incorporate additional references to industry frameworks or guidelines, such as from OWASP or MITRE, or expanding the Guidelines to address additional aspects such as zero-trust and least-privilege architectures, identity management for AI agents, and credentials and secret management for AI.
- 10.35. One respondent sought more guidance on expectations on penetration testing, red teaming and vulnerability assessment for AI systems. Specifically, whether both AI-specific adversarial testing (e.g., jailbreak attempts or prompt injection), and infrastructure-based red teaming (e.g., attacks on APIs or



plugins connected to the AI system) are expected; whether such testing should be manual, automated or combination of both; and whether the expectations extend beyond traditional infrastructure scanning (e.g., container security and CVE scanning) to include AI-specific artifact scanning (e.g., model weights or files for embedded malware or pickle serialisation attacks).

MAS' Response

- 10.36. MAS agrees that controls to manage technology and cybersecurity risks should be applied proportionately to the risk materiality of the AI use case(s) supported by the system. The frequency and rigour of technology and cybersecurity reviews may be calibrated to the risk materiality of the AI use case(s). MAS has incorporated the notion of risk proportionality within the AIRG section on Technology and Cybersecurity Risks.
- 10.37. The Guidelines set out MAS' expectations on technology and cybersecurity risk management in FIs' AI deployments, including controls relating to AI system identities and access, third-party AI governance, and adversarial testing and safeguards for AI systems. These expectations complement MAS' existing expectations under the Technology Risk Management Guidelines (TRMG) and should also be read in conjunction with the MAS Notices on Cyber Hygiene and Technology Risk Management. For example, an AI system that meets the criteria of "critical system" under MAS' Notice on Technology Risk Management will continue to be subject to the availability and recoverability requirements set out there. Where applicable, existing controls implemented by FIs to meet the MAS Notices and Guidelines may also be applied to their AI systems and components.
- 10.38. MAS recognises that an FI may leverage relevant industry frameworks and guidelines to meet the expectations under the Guidelines. MAS has updated the Guidelines to include references to the Cyber Security Agency of Singapore's Guidelines and Companion Guide on Securing AI Systems as resources that FIs may find useful in managing cybersecurity risks in AI systems.
- 10.39. In relation to penetration testing, red teaming and vulnerability assessment, the Guidelines do not prescribe specific testing methodologies or techniques. FIs should adopt a risk-based and proportionate approach, with nature, scope and depth of testing commensurate to the AI system's risk profile, architecture and deployment context. Depending on the use case, this may include AI-specific adversarial testing (e.g., prompt injection and jailbreak attempts), traditional infrastructure security testing (e.g., testing of APIs, containers and dependencies), or a combination of both.
- 10.40. FIs may deploy an appropriate mix of manual and automated testing techniques, as appropriate, to provide reasonable assurance that both known vulnerabilities and emerging AI-specific risks are



adequately identified and addressed, taking into the account the AI system and its supporting environment.

Reproducibility and Auditability

- 10.41. A few respondents requested for greater clarity on how proportionality should be applied in determining the extent of reproducibility and auditability required. Two respondents sought guidance on the extent to which logs should be kept on AI use, such as on prompts and associated outputs of Generative AI, to facilitate auditability.
- 10.42. A few respondents raised concerns with the expectation that documentation be sufficiently detailed to allow for potential replication of the AI system or model and its results. It was highlighted that certain types of AI such as Generative AI were non-deterministic in nature, for which replicating model outputs for such AI would not be achievable in practice even if aspects such as the model version or input were controlled for. Respondents suggested various alternatives to requiring potential replication of results, such as the development methodology, testing process, or stability of the outputs.

MAS' Response

- 10.43. The intent of documenting the AI development process is to support visibility into the risks and limitations associated with an AI use case, and how such risks have been addressed during development. It also facilitates auditability by providing relevant information to parties not directly involved in the development process, such as pre-deployment reviewers or independent validators. The Guidelines have been updated to clarify that the extent and rigour of documentation should be commensurate with the risk materiality of the AI use case.
- 10.44. MAS agrees that expecting documentation of the AI development process to enable replication of exact results may not be an achievable or appropriate goal for non-deterministic AI, such as Generative AI. The Guidelines have been revised to clarify that the documentation process should enable replication of the evaluation and testing process, rather than the exact model outputs, for the purpose of facilitating audits and independent verification. With regards to third-party AI, this refers to the evaluation and testing of the third-party AI when applied on the FI's own data and use cases.
- 10.45. MAS agrees that documentation should not be confined to the development process; clear records should also be kept on post-deployment activities such as actions, decisions, recommendations, and outputs made by, or with the assistance of AI. For Generative AI and AI agents, logging prompts and model responses may also be useful to facilitate ongoing monitoring, human oversight, and detect

unusual activity. Accordingly, paragraph 5.23.(d) of the Guidelines has been amended to indicate that enhanced documentation may be considered for Generative AI or AI agents, such as logging prompts and their outputs. An FI retains flexibility to determine the relevance and necessity for such enhanced documentation, taking into account the risk materiality of the AI use case.

Human Oversight

10.46. A few respondents sought guidance on the level of human oversight expected and degree of autonomy that may be permitted for AI use cases of different risk materialities. One respondent sought further clarification on the minimum standards expected for human oversight and testing, particularly for lower-risk AI systems. Another respondent suggested explicitly disallowing fully automated decisions for high risk use cases.

10.47. One respondent highlighted the importance of decision-deferral mechanisms, where an AI system refrains from making an autonomous decision and escalates the case to human oversight after predefined safety thresholds are breached.

MAS' Response

10.48. The level of human oversight and degree of autonomy should be commensurate with the risk materiality of the AI use case. An FI should determine what constitutes an appropriate level of oversight and degree of autonomy, and be able to demonstrate to MAS how its approach appropriately mitigates risks associated with the AI use case, consistent with the FI's risk appetite.

10.49. Higher-risk AI use cases, particularly those that impact customer outcomes, typically warrant stronger human oversight and lower degrees of autonomy. For example, an FI may deem it appropriate to require human review of all decisions for higher materiality use cases such as credit decisioning or insurance underwriting. Conversely, for lower-risk or assistive use cases, a post-hoc monitoring or review mechanism may suffice.

10.50. The Guidelines have also been updated to recognise decision-deferral mechanisms as one example of how AI systems can be designed to facilitate appropriate human oversight.



Pre-Deployment Reviews

- 10.51. A few respondents raised concerns on the requirement for AI use cases of high risk materiality to undergo formal independent validation before deployment. They noted that the expertise and familiarity needed to assess an AI use case may reside primarily with the development team, and FIs may have limited internal resources with both the requisite expertise and independence to conduct a meaningful validation.
- 10.52. Respondents also sought greater clarity on the degree of independence required, including the functions or units that could conduct formal independent validation. One respondent asked whether reviews commissioned or coordinated by third-party AI providers could qualify as independent validation.
- 10.53. A few respondents sought clarity on the scope of pre-deployment reviews, including whether AI developed and deployed globally would require separate local-level review, and the extent to which conceptual soundness of the design should be evaluated for Generative AI models.

MAS' Response

- 10.54. MAS considers formal independent validation as an important expectation for AI use cases of high risk materiality. This provides objective and effective challenge to the development and deployment process before the FI assumes the associated risks. Independent validation also supports board and senior management oversight by providing assurance that relevant controls have been adhered to, especially in circumstances where failure and poor performance of AI can have a material impact on the FI and its stakeholders. Formal independent validation is not expected in all instances; other forms of documented review may be appropriate for AI use cases not of high risk materiality.
- 10.55. MAS does not prescribe the function or party that must conduct formal independent validation. An FI may determine the appropriate arrangement, including engaging external expertise, provided that the validator meets the expectations in the Guidelines, including having sufficient expertise and independence from the development and deployment teams.
- 10.56. MAS recognises that developers would typically have greater familiarity with the AI use case. However, this does not preclude meaningful independent validation, nor does independence require validators to work in isolation from developers. Validators may engage developers to understand the use case, clarify documentation or replicate results. However, validators should exercise independent judgement and retain responsibility for the validation conclusions. Adequate documentation of the AI



development process, as required under paragraph 5.17. of the Guidelines, would support this assessment.

- 10.57. MAS appreciates that putting in place an appropriately qualified and independent validation arrangement may require additional resourcing. Where an FI is unable to do so for an AI use case of high risk materiality, it should consider whether the risks of the use case can be reduced through changes to its design, scope or deployment. If the use case remains of high risk materiality, the FI should assess whether it is appropriate to proceed with deployment.
- 10.58. Reviews commissioned by a third-party AI provider may qualify as a form of “documented review” under paragraph 5.20. of the Guidelines. However, such reviews would only qualify as a formal independent validation if the FI is able to determine that the review was conducted by a party with the necessary expertise and independence from the development and deployment teams of the third-party AI provider. Regardless of whether an FI relies on such a review, it remains responsible for assessing the suitability of the third-party AI for its intended use. This includes appropriately testing the AI on its own data, and maintaining sufficient documentation of such testing, as set out in paragraph 5.10. of the Guidelines. Where an existing review does not adequately address the applicable validation expectations, the FI should commission or conduct its own review as needed.
- 10.59. AI developed and deployed at global level need not, in all cases, be subject to separate local-level pre-deployment reviews. As stated in paragraph 1.5. of the Guidelines, an FI that is part of a global group may leverage the AI risk management frameworks of its parent entity to meet the expectations in the Guidelines. Consistent with this principle, an FI may rely on pre-deployment reviews conducted at global level, where these meet the relevant expectations in the Guidelines. In particular, the FI should ensure that:
- (a) AI use cases of high risk materiality are subject to formal independent validation, as per paragraph 5.19. of the Guidelines; and
 - (b) the review and testing approach adequately addresses the context in which the AI will be used by the FI, including whether testing datasets are representative of the local context of use, in accordance with paragraph 5.14. of the Guidelines.
- Where the global review and testing approach does not adequately address the local context of use, the FI should conduct additional local-level review or testing to address the identified gaps before deployment.
- 10.60. MAS recognises that the nature of effective validation will differ across AI technologies. The areas listed in paragraph 5.19.(a) to (f) of the Guidelines are therefore not intended to operate as a prescriptive checklist. Validators should assess the areas relevant to the particular AI use case and apply validation



approaches appropriate to the nature of the technology. For example, for Generative AI, assessing conceptual soundness need not entail demonstrating theoretical properties of the underlying model where this is not practicable or meaningful. The Guidelines have been clarified to indicate that validation should cover areas relevant to the AI use case.

Post-Deployment Monitoring

- 10.61. A few respondents sought further guidance on monitoring expectations, including examples of monitoring metrics and thresholds to be used, particularly for dynamic AI or AI which were frequently updated. Respondents also sought clarity on the expected frequency of monitoring. One respondent sought guidance on whether monitoring should be conducted at the use case level, or as an aggregate, risk portfolio analysis.
- 10.62. A few respondents sought further guidance on expectations regarding reviews and re-validations of AI, including the extent of independence required, and expected frequencies of such reviews and re-validations.
- 10.63. A few respondents sought clarity on expectations regarding incident and issue management, including requests for clearer definitions on what constituted an incident, and whether MAS intended to specify severity thresholds or classification frameworks, together with corresponding control or response requirements. Respondents also enquired as to whether MAS envisaged any specific requirement for an FI to report AI incidents to MAS, over and above existing regulatory requirements.
- 10.64. A few respondents sought clarity on MAS' expectations regarding contingency planning and fallback options. In particular, they asked what MAS expected a contingency plan to cover, and whether such plans were required for AI use cases that were not of high risk materiality. One respondent asked if contingency plans should be developed at the FI level, or separately for each use case of high risk materiality.
- 10.65. Respondents also asked for guidance on the circumstances under which MAS expected "kill switches" as opposed to other fallback options, such as rate limiting or scoped disablement. One respondent noted that contingency plans and fallback options are already integral to existing IT and business continuity planning (BCP) frameworks, and that these considerations may best be managed within existing IT and BCP frameworks, rather than requiring separate AI-specific measures in all cases.



MAS' Response

- 10.66. Monitoring activities should be conducted at the use case level, given that risk materialities and applicable monitoring measures could differ between AI use cases. Similar to the section on evaluation and testing, MAS will not prescribe specific monitoring measures and thresholds, given the diversity of AI applications across the industry. In determining appropriate monitoring measures and thresholds, an FI may take reference from industry best practices, and should calibrate these in accordance with the risk materiality of the AI use case. Measures established during pre-deployment review may be used as a starting point, with additional measures incorporated where appropriate to detect performance degradation or drifts, such as data or concept drifts.
- 10.67. As stated in the Guidelines, the frequency and intensity of monitoring should be proportionate to risk materiality. More stringent monitoring would be expected for AI that are designed to be updated automatically, such as adaptive or reinforcement-learning-based models where behaviour may materially evolve during operation without human involvement. An FI has flexibility in determining an appropriate monitoring approach, whether via continuous monitoring or periodic assessments, with frequency informed by the risk materiality of the AI use case.
- 10.68. The degree of independence expected for reviews and re-validations of deployed AI should be similar to pre-deployment reviews. For AI use cases of high risk materiality, re-validations should be performed by competent personnel or functions possessing the necessary expertise and objectivity, and who are independent from the development and deployment teams. For AI use cases not of high risk materiality, other forms of documented review may be performed. The Guidelines have been updated to clarify that the need to conduct re-validations or reviews, as well as the frequency and rigour of such re-validations or reviews, should be informed by risk materiality.
- 10.69. The Guidelines now define “AI incidents” as “incidents that involve a failure, malfunction, or poor performance of the underlying AI system(s) or model(s).”. An FI should establish its own processes to identify, manage and escalate such AI incidents, including severity thresholds or classification frameworks where relevant, to appropriately address such incidents. MAS does not intend to introduce AI-specific incident reporting requirements at this juncture. An FI should continue to comply with existing regulatory requirements on incident reporting to MAS (e.g., on Technology Risk Management or Business Continuity Management), irrespective of whether the root cause of the incident is AI-related.
- 10.70. Contingency plans are expected for AI use cases of high risk materiality. For other AI use cases, an FI has flexibility to determine whether contingency planning is necessary. As not all fallback options are equally viable across AI use cases, contingency plans should be sufficiently specific to the use case. An FI also has flexibility in determining appropriate fallback options, including whether “kill switches”

should be adopted as opposed to other measures. Existing IT and BCP frameworks may be leveraged, provided that they meet the expectations in the Guidelines on contingency planning, including clearly setting out appropriate fallback options for each AI use case of high risk materiality and ensuring that the plans are operationally feasible under a range of operating conditions.

Change Management

- 10.71. One respondent highlighted that the Guidelines appeared to treat re-training of a model as a relatively minor change. However, re-training non-linear machine learning systems could produce logically different models and materially different outcomes.
- 10.72. One respondent sought clarity on whether vendor-driven updates should be subject to the requirement that significant changes trigger appropriate reviews and re-approval processes before implementation. It was highlighted that an FI might not be able to halt vendor development regardless of the FI's internal review and re-approval processes.

MAS' Response

- 10.73. MAS has updated paragraph 5.25. of the Guidelines to require an FI to establish a change assessment framework to evaluate the significance of changes to deployed AI, based on the nature of the change and assessed risk materiality of the AI use case. Based on this framework, an FI may determine for itself the circumstances under which re-training of a model constitutes a significant change, and subject such significant changes to appropriate reviews and re-approvals before implementation.
- 10.74. The change assessment framework mentioned in paragraph 5.25. of the Guidelines should also apply consistently to vendor-driven updates of third-party AI. As mentioned in paragraph 9.14., an FI should also ensure that its contractual arrangements with third-party providers allow the FI to have a risk-proportionate degree of visibility into updates or changes to third-party AI that could reasonably be expected to affect the continued performance, behaviour or risk profile of the AI use case. MAS recognises that vendor-driven updates may occur before an FI is able to complete assessments on the significance of the updates, and subject significant updates to internal reviews and re-approvals. Where such third-party AI relate to use cases of high risk materiality, an FI should implement compensating controls such as enhanced monitoring to mitigate the risks associated with significant vendor-driven updates proceeding without an FI's knowledge or prior assessment.

Requirements for pilots / partial deployments

- 10.75. Two respondents sought clarification on minimum lifecycle controls expected before a pilot/partial deployment (as stated in footnote 28 of the Guidelines) could occur, and what triggers the requirement to complete full pre-deployment independent validation.
- 10.76. One respondent suggested to require pilots to be explicitly time bounded with defined scope and explicit success criteria, and for any extension to require explicit re-approval and a refreshed risk assessment, as part of instilling greater discipline.

MAS' Response

- 10.77. The intent of footnote 28 is to recognise that an FI may find it necessary to deploy an AI use case in a limited fashion in order to establish its suitability. For example, an FI may need to deploy AI to obtain sufficient representative real-world data, for the purposes of evaluation and testing.
- 10.78. Prior to the start of the pilot or partial deployment, an FI should establish appropriate controls and guardrails, guided by the principle that residual risks should remain within the risk appetite of the FI. An FI's policies on the use of pilots or partial deployments should clearly set out such controls and guardrails, which may include time-bound deployment, restricting user access, or other appropriate measures. An FI should also ensure that such controls and guardrails are robustly designed and implemented, given the possibility for highly autonomous AI to devise ways to evade or circumvent these controls.
- 10.79. As set out in paragraph 5.19. of the Guidelines, AI use cases assessed as having high risk materiality should undergo formal independent validation before deployment; this applies irrespective of whether such use cases are pilots or partial deployments.

11. AI Capability and Capacity

- 11.1. MAS had proposed that an FI should determine and ensure the necessary competence and proper conduct of personnel involved in developing and deploying an AI use case, system or model. Regular reviews should be conducted to ensure that the relevant personnel involved in developing, deploying and maintaining an AI use case, system or model are equipped with adequate capabilities and capacity



for effective AI risk management. Technology infrastructure should also be sufficient to meet performance, scalability, and resilience needs for AI.

Scope of expectations

11.2. Two respondents suggested that expectations on ‘competence and proper conduct’ could be cast in a wider context and not constrained to developers and deployers. For example, risks could arise from mis-use of AI by personnel who are not developers or deployers.

11.3. A few respondents sought clarification on the technical qualifications or minimum competencies required of AI developers and deployers. One respondent suggested that expectations on AI training and capacity building be applied in a risk-proportionate manner, and treated as good practices rather than uniform requirements.

MAS’ Response

11.4. MAS agrees that an FI should ensure competence and proper conduct not just for personnel involved in AI development and deployment, but also personnel who are involved in using AI. The Guidelines have been amended accordingly.

11.5. Expectations on AI training and capacity building should be applied in a risk-proportionate manner. The level of technical qualifications and minimum competence should be proportionate to the risk profile of AI used in the FI. An FI with complex AI systems and models should endeavour to have higher levels of technical competence within their development and deployment teams, particularly where such AI will be used in high risk use cases.

11.6. It is important for an FI to possess the necessary capabilities to understand the AI systems and models that underpin its use cases, the risks and limitations, and how to best manage such risks. Hence, an FI should regard AI training and capacity building as supervisory expectations rather than merely good practices.

Regular reviews of adequate capabilities and capacity

11.7. One respondent highlighted that the Guidelines suggested a need for regular, documented formal reviews on the relevant personnel involved in developing, deploying and maintaining an AI use case,



which could be onerous. It was suggested to provide an FI flexibility to decide on the appropriate manner to keep its AI risk management capabilities current.

MAS' Response

11.8. The intent of the review process is to ensure that an FI maintains adequate capabilities and capacity for effective AI risk management on an ongoing basis. For example, an FI which previously focused on traditional AI may need to update its training programmes and materials as it scales up their use of more complex technologies such as AI agents.

11.9. An FI has flexibility to determine an appropriate form and frequency for conducting such reviews. For example, instead of formally assessing the competencies of individual personnel, an FI may consider assessing whether personnel have completed training that is relevant for their functional role, and whether the training materials remain current and appropriate for AI used in the FI.

Technology infrastructure

11.10. One respondent suggested for expectations on technology infrastructure to be made more principle based and technology-neutral, instead of prescribing specific infrastructures. This respondent also suggested to align expectations with the existing Technology Risk Management Guidelines to ensure integration.

MAS' Response

11.11. MAS' expectations on technology infrastructure are intended to be read in a principle based and technology neutral manner. Details of infrastructures such as “graphics processing units, network infrastructure, system memory, secure data pipelines” are illustrative examples, meant to guide an FI in assessing whether its technology infrastructure is adequate for its use of AI.

12. Generative AI and AI Agents

Discussion of risks

12.1. One respondent suggested that the discussion of risks posed by Generative AI in Section 1 of the Guidelines could inadvertently discourage an FI from investigating or implementing available mitigants on the belief that such mitigants would not be acceptable. Another respondent suggested that the Guidelines explicitly set out that there are “no restrictions” on the use of Generative AI, or on the types of data which may be used, so long as risks are properly understood, assessed, and addressed within a FI’s established risk appetite. This respondent also suggested more explicit statements on permissible uses of Generative AI, to provide more certainty on whether an FI may deploy Generative AI solutions.

MAS’ Response

12.2. The discussion of risks posed by Generative AI and AI agents in Section 1 of the Guidelines is intended to alert FIs to how newer AI technologies, specifically Generative AI and AI agents, may amplify risks. The intent is to enhance awareness and understanding of such risks, especially by FIs with less experience deploying such technologies, and not to discourage or restrict their use.

12.3. MAS does not intend to restrict the use of any specific AI technology or limit use cases, provided that the risks are adequately addressed. An FI should determine how best to deploy AI in a safe and responsible manner, taking into account the rigour of its controls, the application of appropriate risk mitigants, and its own appetite for risks from AI use.

12.4. MAS recognises that agentic AI introduces a distinct set of risks that warrant dedicated attention. This is given its capacity for autonomous decision-making, tool use, and chaining of actions across multiple steps or agents. As the understanding of these risks continues to develop, FIs may refer to relevant literature and industry publications to stay abreast of emerging risks and evaluation practices. MAS will update the Guidelines to reflect material developments in the Generative AI landscape and supervisory expectations at a later stage.

Application of controls

- 12.5. A few respondents requested for further guidance on evaluation and testing for Generative AI. In particular, they highlighted practical challenges in evaluating such models, noting that traditional evaluation methodologies may not be fully applicable. Feedback was also provided on sections that could be further adapted for Generative AI, including data management, evaluating and testing methods to assess failure modes such as hallucinations, and explainability considerations. One respondent requested that the Guidelines include minimum controls specific to Generative AI.
- 12.6. A few respondents requested for further guidance on the application of lifecycle controls for agentic AI, including how risks should be evaluated at the level of individual agents, as well as at the system level, taking into account amplification of risks across chains of agents. Additional guidance was requested on risk-proportionate monitoring coverage across the action space. One respondent also highlighted that some agentic AI systems may autonomously retrieve, combine, or synthesise external information, raising distinct considerations such as verifying the provenance of retrieved information, assessing reproducibility of outcomes, and documenting reasoning pathways.

MAS' Response

- 12.7. The Guidelines set out the outcomes and risk management principles that FIs are expected to achieve, affording institutions the flexibility to adopt evaluation approaches that are most suited to their specific Generative AI or agentic AI application. MAS will not mandate specific testing techniques or minimum controls at this time. Nonetheless, MAS has updated relevant sections in the Guidelines to provide further illustrations where appropriate, including examples that are relevant in the context of Generative AI.
- 12.8. With regard to agentic AI, the Guidelines clarify that monitoring should extend, where relevant, to information flows and decision-making paths across workflows that use AI, such as reasoning processes, actions taken, and tools used. Such monitoring can support the identification of individual actions taken across chains of agents, and help pinpoint the sources of specific failures.
- 12.9. Given the pace of developments in this area, MAS is taking an iterative approach. In the near term, MAS has included the national-level Agentic AI Model Governance Framework published by the Infocomm Media Development Authority as a practical reference for managing risks associated with agentic and multi-agent systems. Over time, MAS intends to consult the financial sector separately on what additional guidance specific to agentic AI would be most helpful, taking into account the sector's

evolving use cases and risk experience. MAS welcomes continued engagement from industry on this topic.

13. Implementation Timeline

- 13.1. MAS had proposed a transition period of 12 months after the Guidelines are issued, for FIs to assess and implement the Guidelines.
- 13.2. Many respondents found the proposed transition period challenging, and requested further time to implement the Guidelines, with 18-24 months being the most commonly requested period. Challenges cited included the need to source for sufficient expertise/talent to ensure adequate capabilities and capacity for AI risk management, the need to identify and apply the Guidelines in respect of third-party AI or retrospectively on AI which had already been deployed, or the need to align with Group-level policies on AI risk management.
- 13.3. Some respondents also proposed a phased transition period, whereby certain aspects of the Guidelines would be implemented by 12 months, and additional time taken for full implementation.

MAS' Response

- 13.4. MAS agrees to allow FIs with more time to fully implement the Guidelines in areas which are more resource-intensive.
- 13.5. Unless paragraph 2.3. of the Guidelines applies, an FI should meet the expectations in Sections 3 and 4 of the Guidelines (i.e. establish frameworks for oversight of risks from AI use, and processes for identification, inventorisation and risk materiality assessments) by 12 months. Full implementation of the Guidelines including Sections 5 and 6 (i.e. the application of lifecycle controls across the FI's use cases, and ensuring adequate capabilities and capacity) should be completed by 24 months. Notwithstanding, an FI should apply appropriate lifecycle controls for its high risk use cases as soon as possible, and not delay application until the end of the 24-month transition period.
- 13.6. An FI may apply basic AI governance policies and procedures for the use of AI if they meet the criteria in paragraph 2.3. of the Guidelines, and should complete implementation by 12 months.



Annex A

List of respondents to the consultation paper on Guidelines on Artificial Intelligence Risk Management

1. /TMP Labs Pte. Ltd. (Singapore)
2. Access Partnership
3. AI Ling Advisory Limited
4. Alternative Investment Management Association
5. Andrew Koh M W
6. Aon Singapore Pte Ltd
7. Asia Pacific Loan Market Association Limited
8. Asia Securities Industry & Financial Markets Association
9. Asia Wealth Platform Pte. Ltd.
10. askOdin Pte Ltd
11. AXA XL which includes "Catlin Singapore Pte Ltd", "XL Insurance Company SE (Singapore) Branch" and "XL Re Europe SE (Singapore) Branch"
12. AXS Pte Ltd
13. Behavox Ltd
14. Blitz Consulting & Coaching
15. CFA Society Singapore



16. Chenyi Ang
17. Chris Fong Kar Hoe
18. Clearspring Insights
19. Commerzbank Aktiengesellschaft Singapore Branch
20. CPA Australia
21. Curia Regis Pte Ltd
22. Cyber Risk Institute
23. Cyber Sierra (Fort One Technologies Pte. Ltd.)
24. CyberArk Software Pte Ltd
25. DDN LLC dba Digital Directors Network
26. Diligent Risk Solutions Pte. Ltd.
27. Edwin Hooi
28. Endowus Singapore Pte Ltd
29. Ensign InfoSecurity
30. Ernst & Young
31. Fibish Madathil
32. FIL Investment Management (Singapore) Limited
33. FinbotsAI
34. FIX Trading Community



35. Flagright
36. FWD Singapore Pte Ltd
37. The Global Centre for Risk and Innovation, The Global Risks Forum, The Global Risks Alliance (joint submission)
38. Global Legal Entity Identifier Foundation
39. Grant S Case
40. Holland & Marie Pte Ltd
41. Interactive Brokers Singapore Pte Ltd
42. International Organization for Standardization (ISO), Standards Advisory Group (SAG)
43. Investment Company Institute
44. Investment Management Association of Singapore
45. James Lee
46. Jianjiang Yang
47. Kelvin K. F. Law
48. Lloyd's of London (Asia) Pte Ltd
49. Loh Wai Kwan
50. Marsham Edge Pte Ltd
51. Masato Tsukada
52. MindXO Consulting FZCO
53. MUFG Bank, Ltd.



54. Omnitrade Technologies Pte Ltd
55. Oversea-Chinese Banking Corporation Limited
56. Palo Alto Networks
57. Ping Identity
58. PricewaterhouseCoopers Risk Services Pte Ltd
59. QCP Trading Pte Ltd
60. Rapyd Holdings Pte Ltd
61. RHB Bank Berhad Singapore Branch
62. Sachin Shah
63. Silent Eight Pte. Ltd.
64. Singapore FinTech Association
65. Singapore Reinsurers' Association
66. Skandinaviska Enskilda Banken AB (publ) Singapore Branch
67. SonarSource Sàrl
68. Sumitomo Mitsui Banking Corporation
69. Tanveer Bhatti
70. The Hongkong and Shanghai Banking Corporation (Singapore Branch), HSBC Bank (Singapore) Ltd
71. Thunes
72. Transamerica Life Bermuda



- 73. Transparently Pte Ltd
- 74. UAPF / Algomation
- 75. United Overseas Bank Ltd
- 76. UOB Kay Hian Private Limited
- 77. US-ASEAN Business Council
- 78. Utmost International Isle of Man Limited, Singapore Branch
- 79. Utmost Worldwide, Singapore Branch
- 80. Valli AI FinTech Cybersecurity Technologies LLC
- 81. Veraxis Research Group
- 82. VeritasChain Standards Organization (VSO)
- 83. Vinay Batnur
- 84. Wiz, Inc.

Please refer to Annex B for the submissions.

Respondents who requested for confidentiality for their submission

- 85. AIA Singapore Private Limited
- 86. Amazon Web Services
- 87. Ceffu SG Pte. Ltd.
- 88. Deloitte Singapore SR&T Pte. Ltd.
- 89. GXS Bank Ltd



90. Intercontinental Exchange Inc.

91. IQEQ Regulatory Compliance (Singapore) Pte. Ltd.

92. MSIG Insurance (Singapore) Pte Ltd

93. SailPoint Technologies Pte. Ltd.

94. Securities Association of Singapore

95. Xtransfer Pte Ltd

Other respondents

96. 11 respondents requested for confidentiality of identity

97. 29 respondents requested for confidentiality of identity and submission