



Monetary Authority of Singapore

Guidelines

7 October 2026

Guidelines on Artificial Intelligence Risk Management

Contents

1.	Introduction	3
2.	Application of the Guidelines	8
3.	AI Oversight	10
4.	Key AI Risk Management Systems, Policies and Procedures	13
5.	AI Life Cycle Controls	17
6.	AI Capability & Capacity	30

1. Introduction

- 1.1. The Guidelines on Artificial Intelligence (AI) Risk Management (the “Guidelines”) set out MAS’ supervisory expectations relating to AI risk management in financial institutions (FIs)¹, and seek to support their responsible adoption of AI. These Guidelines set out MAS’ supervisory expectations in the areas of oversight, key AI risk management systems, policies, procedures, AI life cycle controls, and capabilities and capacity for the use of AI. An overview of the key sections covered in these Guidelines is set out below.




Figure 1: Overview of the Guidelines

- 1.2. The Guidelines complement existing legislation and guidelines issued by the MAS. They should be read in conjunction with relevant guidelines, information papers and circulars that MAS may issue from time to time. For instance, existing principles on Fairness, Ethics, Accountability and Transparency (FEAT)² continue to apply for guiding the use of AI in the financial sector. FIs should also take into account applicable industry standards where appropriate³. The Guidelines apply on a group basis for locally

¹ As defined in Section 2 of the Financial Services and Markets Act 2022.

² <https://www.mas.gov.sg/publications/monographs-or-information-paper/2018/feat>

³ FIs may also refer to key initiatives at the national level, such as the Cyber Security Agency of Singapore’s Guidelines and Companion Guide on Securing AI Systems, the Infocomm Media Development Authority’s Model (IMDA) AI Governance Framework, as well as the various initiatives under the AI Verify Foundation to inform their use of AI.

incorporated FIs, that are (i) subject to consolidated supervision by MAS⁴ or (ii) an owner of a critical information infrastructure⁵.

- 1.3. For the purpose of these 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 deployment⁶. Calculators or tools whose outputs are solely based on predefined programming logic or rules would not be regarded as AI for the purpose of these Guidelines.
- 1.4. In addition, the following definitions⁷ are used:
 - (a) A model is a method or approach which converts assumptions and input data into outputs such as estimates, decisions, or recommendations.
 - (b) A system comprises one or more models and other machine-based components.
 - (c) A use case refers to a specific real-world context that the model or system is applied to.
- 1.5. The Guidelines establish a set of high-level expectations that all FIs should adhere to as they deploy AI, including Generative AI and newer developments such as AI agents. An FI which is part of a global group may take guidance from or leverage their Group's AI risk management frameworks, as long as those frameworks meet the expectations set out in these Guidelines.
- 1.6. FIs should implement these Guidelines in a manner commensurate with the size and nature of their activities, and the extent to which their use of AI could lead to material risks to them. AI can be applied to a wide range of use cases, and the risks associated with different use of AI may vary based on the scale, scope and business models of FIs.

⁴ These include locally-incorporated FIs from the banking and insurance sectors.

⁵ As defined in the Cybersecurity Act 2018.

⁶ The following are illustrative examples of AI techniques and applications that would generally fall within the definition above. This list is non-exhaustive, and newer techniques and applications not listed here should be assessed against this definition. Current examples include machine learning (including supervised, unsupervised, and reinforcement learning), deep learning, natural language processing, computer vision, generative AI (including large language models), and AI agents (including multi-agent systems). For the avoidance of doubt, rule-based systems, tools whose outputs are fully determined by explicit formulae that were manually specified, rather than derived from data, and robotic process automation that follows fixed programmed instructions without learning or adaptation would not ordinarily be considered AI for the purpose of these Guidelines.

⁷ Where a point pertains specifically to an AI use case, AI system or AI model, we will use the respective terms explicitly in these Guidelines.

- 1.7. Given the pace of AI developments, MAS also expects FIs to regularly review the adequacy of their AI risk management efforts against AI developments and address new or accentuated risks that may arise due to such developments.
- 1.8. These Guidelines take effect on 7 October 2027. FIs may meet the expectations set out in Sections 3 to 4 from 7 October 2027, and Section 5 and 6 of the Guidelines by 7 October 2028.

Risks

- 1.9. The use of AI can improve performance across business and functional areas but its complexity and probabilistic nature can lead to greater uncertainty, as well as unexpected or more biased behaviour that is harder to identify compared to the use of simpler methods⁸. The greater complexity of AI also leads to challenges in understanding and explaining its outputs that may be inaccurate or biased. Risks from AI use is often an extension of existing risk categories; such risks include:
 - (a) Financial risks, e.g., when used for risk management, the greater uncertainty and unexpected behaviour of AI could lead to poor risk assessments and consequent financial losses.
 - (b) Operational risks, e.g., unexpected behaviour of AI used to automate an FI's operations could lead to operational disruptions or errors in critical processes.
 - (c) Conduct risk, e.g., use of AI can result in biased outputs, leading to unfair treatment of certain customer groups, a bias for certain investment types over others, or conflicts of interests between an FI and its customers.
 - (d) Financial crime risks, e.g., when used to support anti-money laundering efforts, the greater uncertainty and unexpected behaviour of AI could mean that suspicious transactions go undetected.
 - (e) Reputational risks, e.g., customer-facing systems, such as chatbots, could provide incorrect information or make offensive remarks that lead to negative media attention and reputational damage.
- 1.10. With Generative AI, such existing risks associated with AI may be amplified. The greater complexity of Generative AI gives rise to even greater uncertainty and unexpected behaviour compared to AI. The unstructured nature of Generative AI inputs and outputs, and a lack of established techniques in this

⁸ Such uncertainty or behaviour may arise from the noise that is inherent in training data, training data that may not be representative, or when the model encounters scenarios that are not present in its training data. It may also arise due to data or model drifts over time.

area also make it harder to evaluate and test, as well as understand and explain Generative AI's behaviour and outputs. The diverse and often opaque data sources used in Generative AI training, coupled with difficulties in evaluating bias of Generative AI outputs, could also result in decisions that lead to unfair customer outcomes. Generative AI may also lead to a range of other inter-related risks, such as:

- (a) Security risks, e.g., arising from adversarial attacks on Generative AI systems via prompt injection or data poisoning⁹.
- (b) Privacy risks, e.g., arising from the leakage of confidential or customer data due to the use of third-party Generative AI services, or the inappropriate use of customer data in Generative AI without proper consent.
- (c) Legal and intellectual property risks, e.g., arising from the use of Generative AI trained on data that infringes existing copyrights or patents.
- (d) Third-party AI¹⁰ risks, e.g., arising from over-reliance on a few dominant Generative AI providers for mission critical areas, or using open-source models with poor security controls.
- (e) Operational risks, e.g., potential service disruptions due to Generative AI outages and problems arising from inadequate pre-deployment testing.
- (f) Human-factor risks, e.g., use of shadow AI, poor human oversight over Generative AI used in mission critical areas, automation bias or skill degradation due to over-reliance on Generative AI.

1.11. Similarly, the use of newer technologies such as AI agents, which may be granted greater autonomy and access to tools, could further amplify these risks. For example, an AI agent with access to tools could autonomously execute unauthorised or erroneous actions that are not aligned with an FI's business objectives or in a customer's best interests due to a divergence between pre-defined goals and how the AI agent translated such goals into actions¹¹. Compromised AI agents with access to internal systems and external tools could be used to exfiltrate sensitive data or execute malicious commands at scale, amplifying security risks.

⁹ More details on such risks are covered extensively in MAS' Information Paper on Cyber Risks Associated with Generative Artificial Intelligence. The paper can be accessed at <https://www.mas.gov.sg/regulation/circulars/cyber-risks-associated-with-generative-artificial-intelligence>.

¹⁰ For the purpose of these Guidelines, third-party AI refers to AI systems or models, as well as services or tools incorporating or using AI, that are developed, owned, operated or provided by an external party, and which an FI procures, subscribes to, accesses or otherwise relies upon in support of its business activities.

¹¹ FIs may also reference relevant frameworks on specific considerations of risks and governance practices for Agentic AI, such as IMDA's Model AI Governance Framework for Agentic AI. The paper can be accessed at <https://www.imda.gov.sg/-/media/imda/files/about/emerging-tech-and-research/artificial-intelligence/mgf-for-agentic-ai.pdf>.

- 1.12. MAS recognises that AI adoption and its associated risks will continue to evolve alongside advancements in AI technologies. The Guidelines will be updated or augmented as necessary.

2. Application of the Guidelines

- 2.1. The Guidelines apply to all FIs in a proportionate manner, according to their size and risk profile.
- 2.2. Unless paragraph 2.3. applies, an FI should:
 - (a) Establish frameworks, policies and procedures to oversee its use of AI, including clear identification of AI use cases, robust risk materiality assessments and an adequate AI inventory, as set out in Sections 3 and 4 of these Guidelines. The scope of such oversight, policies and procedures may be proportionate to the size and nature of the FI's activities, its use of AI and its risk profile; and
 - (b) Apply relevant life cycle controls, and establish adequate capabilities and technology infrastructure for the use of AI, as set out in Sections 5 and 6 of these Guidelines. Such controls may be calibrated based on their relevance¹² to the AI use cases, systems or models in the FI. Where such aspects are relevant to the AI use case, system or model, the FI may implement them based on the risk materiality of the AI use case.
- 2.3. An FI may apply basic AI governance policies and procedures, as set out in paragraph 2.5., 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. In assessing the potential impact, the FI should consider:
 - (a) the potential financial, operational, regulatory, legal or reputational impact on the FI; and
 - (b) the potential impact on customers or other stakeholders, including in relation to fairness, ethical conduct and consumer protection.
- 2.4. Examples of AI use¹³ that would generally be regarded as meeting the criteria in paragraph 2.3. are:
 - (a) Using AI to assist in drafting, proofreading or rephrasing emails to customers;
 - (b) Using AI to assist in summarising documents or meeting notes for internal reference;

¹² Certain AI life cycle controls, such as those relating to transparency, explainability, and fairness may be more relevant for AI used in areas such as credit scoring, insurance underwriting, provision of financial advisory or fund management services. For AI used in low risk materiality areas to assist humans and not for decision making, such as copilots to assist in writing, AI life cycle controls relating to selection of the AI systems or models, or certain evaluation and testing methods for AI, e.g., stress testing, may be less relevant. However, AI life cycle controls relating to data management, safety and cybersecurity would be relevant and should be applied in a proportionate manner.

¹³ Humans involved in these AI use cases should be reviewing and checking the outputs before using them.

- (c) Using AI to assist in the initial review or analysis of documents for internal reference;
- (d) Using AI to assist in generating formulas, charts or visualisations for internal reference;
- (e) Using AI image generators to assist in designing marketing or internal materials; and
- (f) Using AI-powered chatbots to assist personnel in locating relevant internal resources, such as policies or procedures.

2.5. Basic AI governance policies and procedures should minimally include the following:

- (a) clear accountability for AI oversight (e.g., by designating a member of senior management as responsible for AI oversight);
- (b) permitted and prohibited uses of AI (e.g., prohibitions on inputting confidential, proprietary or client information into public AI tools), together with effective human oversight (e.g., circumstances when mandatory human review and validation of AI outputs prior to use is required);
- (c) maintenance of an approved list of AI tools, and a defined process for requesting approval of new AI tools;
- (d) staff education on the FI's AI-related policies;
- (e) regular compliance checks; and
- (f) periodic and trigger-based review of whether the FI's AI use continues to meet the criteria in paragraph 2.3.

3. AI Oversight

- 3.1. The board of directors and senior management of an FI play critical roles in establishing and overseeing robust frameworks, policies and procedures to support AI risk management across the FI.
- 3.2. **The board and senior management should ensure that existing risk management frameworks, policies, and practices across the organisation¹⁴ adequately identify, assess, and address risks posed by AI.** Board and senior management should maintain effective oversight of risks from AI use, foster the appropriate risk culture for the use of AI, and ensure that the FI's use of AI would not conflict with its ability to meet regulatory requirements or other supervisory expectations¹⁵.
- 3.3. **The board and senior management should ensure consistent standards, clear accountability, and robust coordination across the FI to manage risks from AI use.** While the FI may adopt varying approaches to AI risk management, such as creating a new centralised function to manage all risks from AI use, or managing the incremental risks from AI use via existing risk management functions, the board and senior management should still ensure consistent and coordinated management of risks from AI use across all relevant functions¹⁶.
- 3.4. The board, or a committee¹⁷ delegated by it, is responsible for:
 - (a) approving and regularly reviewing the overall governance approach for AI risk management, including the strategic direction for AI use and key frameworks, structures, policies and procedures designed to assess and manage risks from the FI's use of AI on an ongoing basis;
 - (b) ensuring that risks from AI use are explicitly addressed within the FI's risk appetite framework, including qualitative statements and quantitative measures for such risks as appropriate¹⁸;

¹⁴ Key existing risk management areas where risks from AI use are relevant include but are not limited to model, operational, reputational, data, technology and cybersecurity, third-party, legal and compliance, financial, conduct, and environmental risks.

¹⁵ For example, an FI should continue to adhere to MAS' Guidelines on Fair Dealing even if it leverages AI to deliver products and services.

¹⁶ For example, by establishing structures to coordinate AI risk management in a cross-functional manner, such as cross-functional committees or lines of communication between functions.

¹⁷ For an FI which is part of a global group, the committee could be a governing body or committee beyond local management that is responsible for the oversight of the FI in Singapore, such as a regional or global committee which exercises oversight of the risks from AI use relevant to the FI in Singapore. For an FI which is not part of a global group, this refers to the board of directors or a board-level committee which is delegated by the board of directors.

¹⁸ For example, the FI could establish a qualitative risk appetite statement that articulates the types and level of risks from AI use that the FI is willing to accept, and/or prohibited use categories. Quantitative risk appetite measures could include thresholds on the number or financial impact of AI-related incidents, number of material AI use cases with dependencies on a single provider, or the number of material AI use cases in breach of established performance thresholds.

- (c) setting clear roles and responsibilities for the board and senior management concerning AI risk management oversight;
- (d) ensuring it has an adequate understanding of AI to provide effective oversight and challenge; and
- (e) ensuring that the FI's approach, risk appetite framework, roles and responsibilities, capabilities and culture for risk management of AI use are regularly reviewed to keep pace with newer AI developments, as well as changes in the FI's risk profile and business strategies.

3.5. Senior management is responsible for:

- (a) ensuring the effective implementation of AI-related risk management frameworks, structures, policies and procedures across the FI, consistent with the FI's risk appetite;
- (b) regularly reviewing the effectiveness of the AI-related risk management frameworks, structures, policies and procedures, making appropriate revisions to keep pace with newer AI developments and changes in the FI's risk profile and business strategies, as well as relevant regulatory requirements;
- (c) instituting appropriate strategies and controls to manage risks posed by AI throughout the entire AI lifecycle;
- (d) establishing clear roles and responsibilities for managing risks from AI use across different business lines and functions¹⁹, including control functions responsible for systems and processes relating to AI identification, inventorisation and risk materiality assessment;
- (e) establishing an internal escalation process for managing material risks and exceptions from AI use, such as AI incidents²⁰ or breaches of risk thresholds, and ensuring appropriate and timely actions are taken;
- (f) updating the board on material AI risk issues in a timely manner; and

¹⁹ Such roles and responsibilities may leverage the three lines of defence model if appropriate. For example, first line functions could comprise model developers and users of AI, and may be responsible for areas such as assessing risks from AI use cases or the application of appropriate human oversight. Second line functions could provide effective and independent challenge to first line functions, and may be responsible for areas such as independent validation and overall compliance with applicable regulations and supervisory expectations. The third line (i.e. internal audit) could provide independent assurance on the effectiveness of the AI risk management framework.

²⁰ For the purpose of these Guidelines, AI incidents pertain to incidents that involve a failure, malfunction, or poor performance of the underlying AI system(s) or model(s).

- (g) ensuring the necessary competence of personnel and allocating adequate resources (such as human, technological, financial resources) for effective AI risk management, including appropriate training and capacity building.
- 3.6. Local senior management remains accountable for ensuring that group frameworks, if adopted for Singapore, adequately address local considerations and regulatory requirements specific to Singapore. Local senior management should have adequate visibility into and input on matters affecting Singapore operations, including access to timely and relevant reports and data that inform decisions on how to manage risks from AI use. There should be clear escalation paths for Singapore-specific issues to be surfaced to the responsible parties to ensure timely decision-making and resolution of local concerns. Local senior management should also be able to demonstrate to MAS how they discharge their oversight responsibilities, even when relying on group frameworks.

4. Key AI Risk Management Systems, Policies and Procedures

- 4.1. **An FI should ensure that its AI risk management framework encompasses key systems, policies and procedures for the identification, inventorisation, and risk materiality assessment of AI.** The FI should apply consistent and robust approaches to identify, inventorise and determine the risk materiality of AI, and apply controls proportionate to the assessed risk materiality.

AI Identification

- 4.2. **An FI should establish systems, policies and procedures to ensure the consistent identification of AI use across all relevant business and functional areas.** This identification process is a critical prerequisite for applying the appropriate governance, risk management standards, and controls effectively to AI throughout their lifecycle. Clear AI definitions, criteria and processes, supported by tools and systems as appropriate, should be implemented to facilitate this identification process. Where AI is incorporated within²¹, or used in the delivery of, third-party services, the identification process should minimally cover such AI in services provided by material third-party service providers.
- 4.3. **An FI should assign clear roles and responsibilities for AI identification, including the designation of a control function to be responsible for AI identification systems and processes.** While the actual identification of AI use may be carried out by relevant business or functional units, the designated control function should retain sufficient and independent oversight to ensure the consistent application of the identification process across the FI including the identification of new AI use over time. It should also be the final arbiter on whether a given use constitutes AI use. The designated control function should ensure that clear documentation of the identification process and outcomes is maintained, and that identification systems and processes are regularly reviewed and updated to take into account newer AI technologies or introduction of AI into existing third-party services.
- 4.4. **An FI should ensure that residual risks from unidentified AI use remain within the FI's risk appetite.** Where there are practical constraints in consistently identifying all AI use, e.g., due to inadequate

²¹ For example, software-as-a-service which may not be explicitly sold as AI, but contains AI features that an FI may access or use as part of consuming the service.

disclosure of AI embedded in third-party services or the use of shadow AI, the FI should identify risks that arise from such constraints, and put in place mitigating measures²².

AI Inventory

- 4.5. **An FI should establish and maintain an accurate inventory of AI use cases, systems or models across the FI to the extent possible or practicable, so as to support governance and oversight, as well as risk management, throughout the AI lifecycle.** There should be clear policies and procedures on the maintenance of the inventory, including establishing an appropriate frequency to update the inventory to reflect new, updated or decommissioned AI use cases, systems or models.
- 4.6. **The FI can enhance existing inventories to include AI use cases, systems or models or establish a dedicated AI inventory.** In either case, there should be clear linkages between the AI inventory and other relevant inventories²³ in the FI, to facilitate cross-referencing between inventories as part of effective oversight.
- 4.7. **The AI inventory should capture key attributes at an appropriate level of granularity to enable effective governance and oversight, as well as risk management.** Specific attributes may vary based on the FI's context or the nature of the AI use case, system, or model²⁴, but they may include the following: purpose and description, approved scope of use (e.g., jurisdiction), model type, data used, dependencies, lifecycle status, assigned risk materiality rating, model review status, key roles and responsibilities (e.g., owners, developers), and links to essential documentation. Maintaining this information facilitates risk assessment, monitoring of scope adherence, and consistent application of controls.

²² For example, such measures could include clear guidelines and communication to staff on the allowed and disallowed uses of AI, including embedded AI features and public AI tools (e.g., prohibitions on inputting confidential, proprietary or client information into public AI tools), or uplifting technical controls (e.g., network monitoring, data loss prevention) to detect unauthorised use of AI.

²³ Such linkages would depend on the attributes captured in the AI inventory. For example, the AI inventory could incorporate consistent data asset identifiers linking to existing data inventories, or consistent vendor/service identifiers linking to existing third-party or outsourcing registers.

²⁴ Any differences in attributes captured across the FI's AI use cases, systems, or models should be clearly delineated in the FI's policies and procedures, to ensure consistent application. The omission of certain attributes in the AI inventory should not compromise the effectiveness of the FI's oversight and risk management of AI.

- 4.8. **The design of the inventory should be regularly reviewed to ensure that the attributes captured take into account newer AI technologies** where there may be additional relevant attributes²⁵ or guardrails needed, or additional information relating to third-party AI.
- 4.9. **The FI should assign clear roles and responsibilities for the inventorisation of AI, including the designation of a control function to be responsible for establishing policies and procedures relating to the AI inventory.** While maintenance, updates and regular reviews of the inventory may be carried out by relevant business or functional units, the designated control function should retain sufficient and independent oversight to ensure that the AI inventory remains accurate and complete.

AI Risk Materiality Assessment

- 4.10. **An FI should establish an assessment methodology to evaluate the risk materiality of an AI use case based on the nature of its business and the risks arising from the systems and models used²⁶.** The assessment methodology should be applied consistently to perform risk materiality assessments for each AI use case in the FI. Such assessments are critical for calibrating the risk management approach, ensuring that controls are applied proportionately to the risks. This ensures that AI use cases assessed as having higher risk materiality receive more stringent scrutiny and robust controls, supporting effective oversight and confirming that its use remains within approved boundaries.
- 4.11. **The assessment should take into account both the inherent risk materiality of an AI use case before the appropriate risk management controls are applied, as well as the residual risk materiality after risk management controls are applied.** The FI should ensure that the residual risk materiality of the AI use case meets the FI's risk appetite before deployment. The risk materiality assessment methodology and risk materiality assessments for each AI use case should also be regularly reviewed to ensure their continued relevance and appropriateness.
- 4.12. **The risk materiality assessment of an AI use case should consider various risk dimensions relevant to the FI's context,** minimally covering:
- (a) **Impact:** The potential consequences of a failure, malfunction, or poor performance of the AI on the FI (e.g., financial, operational, regulatory, reputational) and its customers or other

²⁵ For example, for AI agents, relevant attributes may include agent-specific identifiers, information on the tools and systems the agent can access, its components, and the guardrails imposed to manage risk.

²⁶ Dependencies between AI systems or models that support or interact with the AI use case should also be considered, where such dependencies could affect the overall risk profile of the AI use case.

stakeholders (e.g., fairness, ethical breaches, consumer protection). The nature and sensitivity of the data processed by the AI system or model should also be considered.

- (b) **Complexity:** Arising from the nature of the AI technology used, the novelty of its application, the data it uses, or the explainability of its outputs. As the understanding of AI technologies evolves across time, this risk dimension could also evolve, e.g., with more research and greater understanding, the complexity of a new AI technology that was not as well understood initially may change. For third-party AI, the level of visibility into the nature of the AI technology or the data used should also be considered.
- (c) **Reliance:** The level of reliance on AI (including third-party AI) for the associated decision or output, including the extent of autonomy granted to the AI system or model and the degree of human involvement or oversight in the process it supports.

4.13. **An FI should assign clear roles and responsibilities for AI risk materiality assessment.** A control function should be designated to establish the risk materiality assessment framework, and to arbitrate or approve the risk materiality of an AI use case. While specific aspects of the assessment process may be implemented by relevant business or functional units, the designated control function should retain sufficient and independent oversight to ensure the consistent application of the assessment process across the FI. This includes being the arbitrator or approver of risk materiality assessments, and ensuring that such assessments are clearly documented and regularly reviewed for accuracy and relevance.

5. AI Life Cycle²⁷ Controls

- 5.1. **An FI should plan for and implement robust controls covering the entire life cycle of an AI use case, system or model and assign clear roles and responsibilities for such controls across different business and functional units.** AI life cycle controls should be regularly reviewed to take into account the use of newer AI technologies.
- 5.2. **An FI may implement the relevant AI lifecycle controls in a proportionate manner, ensuring that they are appropriate to the assessed risk materiality of the AI use case²⁸.** Where there are practical constraints, e.g., where there may not be adequate disclosure of information by the third-party AI provider, the FI should identify risks that arise from such constraints, and put in place the necessary mitigants, such as limiting the use of AI. After putting in place such mitigants, the FI should ensure that residual risks remain within the risk appetite of the FI.
- 5.3. **For an AI use case assessed as high risk, an FI should develop and implement contingency plans outlining fallback options, such as alternative systems or manual processes, to ensure business continuity in case of AI failure or unexpected behaviour.** Plans should be regularly reviewed and tested to ensure that they are operationally feasible under a range of operating conditions. For AI with "kill switches"²⁹, clear contingency activation protocols should be in place and tested regularly.

Data Management

- 5.4. **An FI should put in place data management controls to ensure data used across the AI life cycle is fit for purpose and representative, of high quality, and subject to robust data governance.** General data governance and management standards, in areas such as data ownership, access controls, and intellectual property rights, should be applied to data used for AI, and, where necessary, uplifted to address AI-specific requirements. The FI should ensure existing data governance controls are adequate

²⁷ The AI life cycle refers to its evolution from inception to retirement or decommissioning (adapted from ISO/IEC 22989 – AI Concepts and Terminology).

²⁸ In certain cases, AI may not be fully deployed from the onset and could be partially deployed, e.g., as pilots or phased rollouts. Such partial deployments of AI may require adjustments to lifecycle controls designed for full deployments, e.g., certain testing may only be conducted after information from the partial deployments are available, or pre-deployment reviews may not have been fully completed. For such partial deployments, the FI should establish clear policies and procedures to govern any deviations from standard lifecycle controls. Such policies and procedures may include time and user limits, clear criteria for success, conditions on the terms of use for owners and end-users, close monitoring of usage patterns and outputs for anomalies and ensuring compliance with the limited scope of use.

²⁹ "Kill switches" are usually used to refer to mechanisms that can be used to deactivate AI expeditiously if they exceed risk tolerances.

to address risks arising from the variety and volume of data involved. Key areas that the FI should consider include:

- (a) **Fit for Purpose:** Suitability of data used in an AI use case, system or model based on the intended objective and context, including assessing the use of data against other considerations, such as fairness.
- (b) **Representativeness of Data:** Representativeness of data used for training and testing the AI use case, system or model across a range of real-world conditions, including stressed conditions, under which the AI will be used.
- (c) **Data Quality:** Adequacy of the quality of data used in an AI use case, system or model, including assessing data relevance, accuracy, completeness, and recency; as well as regular monitoring of data quality and checks for anomalies, drifts, and potential bias.
- (d) **Data Classification:** Appropriate data classification processes to guide appropriate use of data in an AI use case, system or model, taking into account criticality and sensitivity of data used.
- (e) **Data Security:** Secure data lifecycle management practices, including timely destruction or sanitisation of training data, model artefacts and outputs once decommissioned or no longer required; appropriate data protection measures, such as data encryption in transit and at rest, secure data handling, including securing the inputs and outputs of the AI use case, system or model.
- (f) **Data Privacy:** Appropriate data privacy measures based on relevant regulatory requirements and guidance relating to data privacy³⁰, such as the need to obtain consent when using sensitive personal data of customers or employees to train the AI models, or processing such data using AI.
- (g) **Auditability and Lineage of Data:** Appropriate documentation of key data management aspects, such as data sourcing, selection, processing and lineage; how data was assessed as fit-for-purpose and representative; and recording approvals and remedial actions related to data.

³⁰ Such as the Personal Data Protection Commission's Advisory Guidelines on use of Personal Data in AI Recommendation and Decision Systems (<https://www.pdpc.gov.sg/organisations/regulations-decisions/regulatory-guidance/advisory-guidelines-on-use-of-personal-data-in-ai-recommendation-and-decision-systems>) and Advisory Guidelines on use of Personal Data in Generative AI (<https://www.pdpc.gov.sg/organisations/regulations-decisions/regulatory-guidance/advisory-guidelines-on-use-of-personal-data-in-generative-ai>).

Transparency and Explainability

- 5.5. **An FI should establish relevant controls³¹ to ensure an appropriate level of transparency and explainability³² for an AI use case.** Transparency is important to support accountability and trust in the use of AI, and enables customers, internal users and other stakeholders to make informed assessments of the risks, reliability and limitations of an AI use case, system or model. It also allows an FI to demonstrate that their use of AI is in line with their stated objectives and risk management standards, and helps to strengthen trust among customers and other stakeholders.
- 5.6. **The degree of transparency and explainability required should reflect the context and nature of the AI use case³³ and its assessed risk materiality.** For example, AI that is heavily relied upon for credit decisioning, insurance underwriting or other regulated activities with high impact on customer outcomes (e.g., provision of financial advisory services or fund management), will require a higher level of transparency and explainability than AI used solely to assist internal users in low-risk tasks. For AI use cases requiring higher levels of transparency and explainability, the FI should pay greater attention to the different features or attributes in data used as inputs and justifications for their use, the ability for users to identify key drivers of the output, the need to inform customers of the use of AI, the consequences of AI-driven decisions, and channels for redress. The needs of the intended audience should also be considered when determining an appropriate manner to provide transparency and explainability for the AI use case³⁴.

³¹ For example, based on areas relating to transparency and explainability covered in the NIST AI Risk Management Framework at <https://doi.org/10.6028/NIST.AI.100-1>.

³² Transparency refers to disclosures made to individuals or groups of individuals related to the use of AI that affects them, as well as the provision of explanations where relevant and requested. Explainability relates to methods used for facilitating an understanding of AI generated outputs or decisions, such as building explainability into AI model or system development, or post-hoc explainability methods.

³³ Such as the extent to which AI-driven decisions are likely to require explanations (e.g., to facilitate decision making by end-users, or to account to the FI's customers) for the AI use case.

³⁴ For example, information disclosed to the FI's customers need not be overly technical or detailed in nature, and could focus on areas such as how the use of AI may affect them, the basis for decisions made, and channels for further enquiry or redress. In contrast, more technical information on model behaviour and limitations may be relevant for internal stakeholders involved in AI evaluation or oversight. FIs may reference relevant frameworks on transparency practices, such as 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>.

Fairness

- 5.7. **An FI should define what it considers “fair” outcomes³⁵ and have appropriate controls to identify and mitigate harmful biases and discriminatory outcomes across the AI life cycle. Controls should be calibrated to the context and nature of the AI use case and its assessed risk materiality.** For example, greater attention should be paid to AI used in areas such as credit decisioning or insurance underwriting that may lead to unfair access or denial of financial services or products offered to individuals.
- 5.8. Where fairness considerations are relevant, the FI should conduct fairness assessments, which may involve defining relevant protected attributes³⁶, assessing whether the use of AI leads to systematic disadvantages for specific groups using appropriate fairness metrics, and documenting the results and any mitigation steps taken.

Human Oversight

- 5.9. **An FI should put in place and regularly review controls to ensure appropriate human oversight over an AI use case, system or model³⁷ across its life cycle.** The need for and degree of human oversight should take into account the objective of using AI, and be proportionate to the risk materiality of the AI use case. An FI should also take into account automation bias and decision fatigue as the use of AI increases in speed and scale. Key areas that the FI should consider include:
- (a) **Roles and Responsibilities:** Clear assignment of roles and responsibilities for human oversight, including escalation and decision-making processes relating to human oversight.
 - (b) **Capabilities:** Equipping competent personnel assigned to monitor AI use with the necessary capabilities, including the necessary authority and ability to intervene.
 - (c) **Design:** Designing and developing AI systems or models from the outset to enable and facilitate appropriate human oversight. These include mechanisms to escalate AI outputs or recommendations for human review where pre-defined conditions (e.g., on reliability and accuracy) are met.

³⁵ The consideration of “fair” outcomes may take reference from relevant frameworks and principles, such as MAS’ FEAT principles (<https://www.mas.gov.sg/publications/monographs-or-information-paper/2018/feat>).

³⁶ Attributes that are substitutes for or highly correlated to the protected attributes should also be considered.

³⁷ This would include all instances where the FI has assessed that human involvement is necessary, regardless of how human oversight is implemented – e.g., human in the loop or over the loop.

- (d) **Documentation and Review:** Establishing processes to document logs and other records generated through development, deployment, operation, monitoring and incident management activities, as well as regularly reviewing human oversight decisions and interventions (including incidents as well as near misses) to assess the effectiveness of human oversight.

Third-Party AI Management

- 5.10. **An FI should ensure that onboarding, development³⁸ and deployment controls for third-party AI are adequate for the risk materiality of the use case involving third-party AI³⁹.** This would include ensuring that contractual agreements with third-party providers allow the FI to have a risk-appropriate degree of visibility over the introduction of AI and subsequent updates or changes, so as to facilitate the application of relevant controls. The FI should ensure that AI developed by third parties are subject to appropriate reviews, which include testing third-party AI services in the context of the FI's use cases (including using the FI's own data), and performing compensatory testing to address informational gaps arising from inadequate disclosures by third-party AI providers. The FI should maintain sufficient documentation on how it has tested the third-party AI and deemed it suitable for use.
- 5.11. **The decision to onboard and use third-party AI is ultimately a decision of the FI, and the FI retains primary accountability for its use of third-party AI, including the possible impact on its stakeholders and on the FI's ability to comply with relevant regulatory requirements.** In assessing whether to onboard or use third-party AI, the FI should consider the following areas, recognising that each consideration may be met to varying degrees depending on the nature and risk materiality of the AI use case. The FI should ensure that the residual risk materiality of the use case which uses or depends on third-party AI is within the FI's risk appetite; where this cannot be achieved, the FI should consider limiting or suspending the third-party AI service, or replacing the relevant third-party provider:
 - (a) **Third-Party AI Transparency:** The extent to which transparency from third-party AI providers (including disclosures or available documentation) enables an FI to assess whether appropriate lifecycle controls have been applied during the development and deployment of such third-party AI – including those relating to data, model, technology and cybersecurity, fairness, and explainability. To address limitations in transparency, an FI may consider reviewing available certifications or external assessments of the AI provider performed by independent and competent parties (i.e. not self-attestations) that provide coverage of key risks from AI use.

³⁸ Such controls would apply where the FI makes modifications to or fine-tunes third-party AI before deployment.

³⁹ MAS expectations and requirements on managing the risks from outsourcing and use of third-party services would apply to the use of third-party AI.

Compensating controls such as additional testing to understand the behaviour of third-party AI, or the application of greater human oversight, may also be employed to address such limitations.

- (b) **Supply Chain Assessments:** The extent to which key third-party and open-source AI models, datasets, and dependencies have undergone supply chain risk assessments and validation, including reviews of model provenance, training data integrity, and other known vulnerabilities.
- (c) **Concentration Risks:** The extent to which the use of third-party AI may contribute to potential concentration risks arising from over-reliance on key third-party AI providers (whether direct or indirect).
- (d) **Change Management:** The extent to which the FI is able to receive notifications of updates or changes to third-party AI post-deployment, and manage and assess the impact of such updates or changes.
- (e) **Contingency Plans:** Whether robust contingency plans are in place to address potential failures, unexpected behaviour of third-party AI, or discontinuing of support by vendors, particularly for third-party AI used in use cases where the risk materiality is high.
- (f) **Legal Agreements:** The extent to which legal agreements set out clear expectations and responsibilities on the AI provider, such as clauses pertaining to performance guarantees, data protection, the right to audit, notification when AI is introduced and updated, or seeking the FI's agreement before incorporating AI.
- (g) **Capabilities:** The level of awareness and capabilities of staff involved in procurement, development, deployment and use of third-party AI.
- (h) **Complexity:** The need to conduct more detailed assessments when using more complex third-party AI services that the FI may have less experience with⁴⁰.

Selection

5.12. **When selecting AI systems, models⁴¹ or features⁴² to use, an FI should consider the objectives and risks of the AI use case.** The FI should document their considerations during the selection process in a

⁴⁰ For example, the FI may need to do more detailed evaluation and testing, as well as technology and cybersecurity checks if they have not had any experience with developing and deploying AI agents provided by third-party providers.

⁴¹ This may include specific AI models (e.g., specific machine or deep learning models) or techniques (e.g., optimisation or finetuning techniques).

⁴² Features refer to the attributes of data points in a dataset, e.g., for data relating to a loan, the income of the obligor and outstanding value of the loan are two possible attributes or features. Feature engineering refers to the process of selecting,

manner commensurate with the risk materiality of the associated AI use case, particularly when selecting more complex systems or models, or less understood features over simpler or conventional alternatives. This may involve balancing factors like the needs of the use case and performance requirements against complexity, the need for fairness, or transparency and explainability. Where relevant, the selection process may also draw reference from research, or accepted industry practice. The FI should also consider incorporating reviews by domain experts or users (e.g., by subject matter experts or users in the relevant business lines or functional units), to ensure that the selection of systems, models or features aligns with the context of the AI use.

- 5.13. Where newer and more complex AI systems or models which are less understood are selected, an FI should carefully weigh the benefits of deploying such AI systems or models against new or heightened risks to the FI, such as hallucination, opaqueness, security risks, and the FI's capabilities to mitigate such risks.

Evaluation and Testing

- 5.14. **An FI should conduct relevant evaluation and testing that is proportionate to the risk materiality of the AI use case.** Each system or model should be evaluated and tested to meet an appropriate level of reliability and safety based on the risk materiality of the applicable AI use case before deployment. An FI should identify key risks arising from the AI use case and set clear, measurable thresholds. Identification of risks and reliability and safety assessments should take into account the context of AI use, and assessments should take reference from best practices⁴³. Evaluation and testing should assess performance under a range of plausible conditions, from real-world scenarios to edge cases. Key areas that the FI should consider include:

- (a) **Evaluation Measures:** Define appropriate evaluation measures aligned with the AI's objectives and establish acceptable performance thresholds for these evaluation measures. Performance thresholds should be clearly defined, documented, and mutually agreed upon by business owners, developers and reviewers.
- (b) **Testing Approaches:** Employ relevant testing methods, such as out-of-sample or out-of-time testing, sensitivity analysis, stability analysis across different data distributions or time periods,

modifying or creating new features from the original attributes of a dataset to improve an AI model's performance, e.g., normalising income of the obligor and outstanding value of the loan to a common scale ranging from 0 to 1; or creating new derived features, such as a debt-to-income ratio, from existing attributes.

⁴³ References for reliability and safety assessments for Generative AI include IMDA's Starter Kit for Testing LLM-Based Applications for Safety and Reliability which can be accessed at <https://www.imda.gov.sg/-/media/imda/files/about/emerging-tech-and-research/artificial-intelligence/large-language-model-starter-kit.pdf>. The Starter Kit is a set of voluntary guidelines that coalesces emerging best practices and methodologies for the testing of LLM-based applications.

sub-population analysis, stress testing (including edge cases and adversarial testing where appropriate), error analysis, and benchmarking against alternatives. Datasets used for testing should be representative of the context of use in the FI.

- (c) **Overfitting Mitigation:** Where possible, implement techniques to prevent overfitting, especially for more complex AI. This may involve favouring simpler models unless higher complexity is justified by a clear performance uplift, applying regularisation techniques⁴⁴, appropriate feature selection, and using robust validation techniques like cross-validation.

5.15. **Where there are risks and limitations associated with AI identified during development, an FI should put in place appropriate controls and guardrails⁴⁵ to mitigate these risks and limitations before deployment.** Evaluation and testing of newer AI developments such as Generative AI and AI agents should cover their key failure modes and the effectiveness of associated guardrails⁴⁶.

Technology and Cybersecurity Risks

5.16. **An FI should ensure that the AI system is secure, well-governed, and supported by appropriate controls to manage technology and cybersecurity risks commensurate with the risk materiality of the AI use case(s) supported by the system.** Relevant regulatory requirements and guidance relating to technology risk management would apply⁴⁷. Key areas that the FI should consider include:

- (a) **Security:** Deployment should occur in secured IT environments, including hardened configurations, network segmentation, and technical controls such as input validation, API authentication, encryption, and data loss prevention.
- (b) **Access Control:** Access to AI components and infrastructure during deployment should be tightly controlled, with role-based access and multi-factor authentication, and robust privileged account access management. Separation of duties, such as assigning different teams to train and test the AI systems, should also be considered.

⁴⁴ Such techniques generally try to limit the number of parameters used so that the trained model is less complex, e.g., some regularisation techniques force less important parameters to values of zero.

⁴⁵ This may include deploying AI as part of pilots or phased rollouts (with appropriate guardrails) as mentioned in footnote 28 where relevant, e.g., to obtain more representative data to mitigate limitations in testing data.

⁴⁶ For example, for Generative AI, evaluation and testing should pay greater attention to hallucinations, generation of undesirable content such as toxic, harmful or biased content, bias, data leakages or disclosures, and vulnerability to adversarial attacks. For more information on the testing of such risks, please refer to IMDA's Starter Kit for Testing LLM-Based Applications for Safety and Reliability.

⁴⁷ Such as MAS' Guidelines on Risk Management Practices for Technology Risk which can be accessed at <https://www.mas.gov.sg/regulation/guidelines/technology-risk-management-guidelines>.

- (c) **Third-party:** Where third-party components or services are used, such as plugins or APIs, controls should be implemented to govern use, restrict data exposure, and monitor for security, compliance or operational risks, such as cybersecurity issues or service disruptions at third-party providers.

Reproducibility and Auditability

5.17. **An FI should document the AI development process to enable reproducibility and auditability.** The extent and rigour of documentation should be commensurate with the risk materiality of the AI use case, and be sufficiently detailed for an independent party, such as a reviewer or auditor, to understand and potentially replicate the implementation of the AI system or model and its evaluation and testing process. Documentation standards should cover the entire development process and may include information such as:

- (a) data sources, processing, and quality checks;
- (b) selection rationale;
- (c) training procedures including code versions, environments, and hyperparameters;
- (d) evaluation measures and performance thresholds, testing approaches, and results;
- (e) explainability analysis, fairness assessments; and
- (f) key assumptions, limitations and mitigants.

Pre-Deployment Reviews

5.18. **Prior to deployment, an FI should subject the AI use case, including its underlying systems or models, to reviews by parties not involved in its development to ensure that the relevant controls, such as evaluation and testing, have been adhered to.** Such reviews should be supported by clear policies and procedures, and assess the suitability, robustness, and performance of AI for its intended purpose within defined boundaries. Roles and responsibilities relating to such reviews should be clearly assigned, and the scope and independence of these pre-deployment reviews should be proportionate to the AI's assessed risk materiality.

5.19. **The scope of these pre-deployment reviews, as well as the degree of independence of the parties undertaking the assessment and reviews, should be proportionate to the assessed risk materiality of**

the AI use case. AI use cases assessed as having high risk materiality should undergo formal independent validation before deployment. Such validations should be conducted by competent personnel or functions possessing the necessary expertise and objectivity, and who are independent from the development and deployment teams. The validation process should provide effective challenge to developers and should cover areas relevant to the AI use case, such as:

- (a) conceptual soundness of the design;
- (b) suitability and quality of data inputs;
- (c) integrity of the implementation;
- (d) evaluation measures, performance thresholds, testing approaches and results;
- (e) explainability analysis and fairness assessments; and
- (f) assumptions, limitations and mitigants.

5.20. For AI use cases not assessed as having high risk materiality, other forms of documented review (such as peer reviews by qualified individuals that were not involved in development or deployment⁴⁸) on key aspects of the development and deployment processes may be undertaken.

5.21. **Review findings, including any identified limitations, necessary remediation actions, or conditions for use, should be reported by reviewers to the relevant approval body.** The relevant approval body should ensure review recommendations are actioned appropriately.

5.22. **An FI should conduct technology and cybersecurity reviews to ensure that AI can be deployed into the production environment in a controlled and secure manner.** Such reviews should cover areas such as technical implementation, system and network security, system resilience, recoverability and operational readiness of AI for deployment. Key areas that the FI should consider include:

- (a) **Secure Design and Access Control:** Review the AI system architecture, including APIs, plugins and third-party services, for secure design, environment segregation, and access controls based on the principle of least privilege. Prior to go-live, complete deployment checklists to confirm that critical controls, including encryption, data loss prevention, firewalls, access restrictions and logging, are operational and correctly configured.
- (b) **Testing and Threat Mitigation:** Conduct vulnerability assessments, penetration testing, and red teaming across AI systems to identify and remediate security weaknesses. Test AI systems against

⁴⁸ This includes both individuals involved in initial deployment of AI, as well as ongoing deployment of AI, e.g., operating the AI system.

adversarial threats (e.g., manipulation, evasion, poisoning) and implement safeguards, such as input validation, throttling and anomaly detection.

Post-Deployment Monitoring and Review

5.23. **An FI should develop and implement comprehensive and robust controls for the ongoing monitoring of all deployed AI (including third-party AI used in the FI).** Given the uncertainties associated with AI, their dynamic nature and the potential for model staleness and performance degradation due to data or model drifts over time, ongoing monitoring is critical to ensure that deployed AI operates as intended and remains fit for purpose over time. The frequency and intensity of monitoring activities should be proportionate to the assessed risk materiality of the deployed AI. Key areas that the FI should consider include:

- (a) **Monitoring Measures:** Based on the risks of the AI use case, define key metrics to be monitored and acceptable performance thresholds for each metric. Metrics may include relevant dimensions like robustness, stability, data quality, and fairness. Tiered thresholds (e.g., early warning threshold levels) should be considered to pre-empt deterioration and trigger early actions. Appropriate checks for data drifts (changes in input data distributions), concept drifts (changes in relationships between inputs and outputs), and overall model drifts should be implemented. Where relevant, information flow and decision-making paths across workflows that use AI, such as reasoning processes, actions taken, tools used should also be monitored.
- (b) **AI Incident and Issue Management:** Establish robust processes for reporting, tracking, escalating, and resolving issues or incidents if breaches or anomalies arise from the monitoring process. Resolution approaches may include retraining, adjustments, redevelopment, or decommissioning of the underlying AI model or system, depending on the severity and nature of the issue. For high risk materiality AI, consider implementing "kill switches" or override mechanisms to rapidly deactivate the AI system if it exceeds risk tolerances. Provide users of deployed AI with appropriate mechanisms to provide feedback and report issues to support continuous improvement, redress and resolution of issues.
- (c) **Roles and Responsibilities:** Assign clear roles and responsibility for ongoing monitoring and managing of incidents of deployed AI, including the designation of an appropriate accountable person.
- (d) **Documentation:** Maintain clear records of monitoring activities, results, identified issues or incidents, and subsequent remediation actions taken for auditability and ongoing risk management. Access to AI systems or models, training data, pipelines, and configuration files should be strictly controlled and logged. Enhanced documentation for Generative AI or AI agents

could also be considered where relevant, such as logging prompts and model responses, together with applicable model versions and reasoning processes used.

- (e) **Training and Awareness:** Equip persons responsible for monitoring with appropriate training to support effective monitoring and incident management; and equip users of deployed AI with the necessary training and awareness to report unintended behaviour of deployed AI.

5.24. An FI should conduct re-validations or reviews of AI in its portfolio where necessary, to ensure that AI continues to be used appropriately and that associated risks remain appropriately managed. The need for re-validations or reviews, as well as the frequency and rigour of re-validations or reviews, should be informed by the risk materiality of the AI use case. AI use cases assessed as having high risk materiality should be subject to regular re-validations by independent parties. Re-validations or reviews may also be triggered by issues or findings on interconnected AI systems or models, alerts or breaches from ongoing monitoring, significant changes to the AI or its operating environment, or the identification of new risks in the external environment, such as regulatory or technological developments.

Change Management

5.25. An FI should develop and implement comprehensive and robust controls for managing changes to deployed AI. Such controls are essential to ensure that any changes do not lead to unintended behaviour, performance degradation or misalignment with the intended use of deployed AI. Key areas that the FI should consider include:

- (a) **Scope of Change:** Establish a framework to evaluate the significance of changes to deployed AI, based on the nature of the change⁴⁹ and assessed risk materiality of the AI use case. Significant changes should trigger appropriate reviews and re-approval processes before implementation. Compensating controls, such as enhanced monitoring, should be applied in relation to high materiality use cases where significant changes may occur without prior review.
- (b) **Change Control:** Implement change control mechanisms (with appropriate human oversight) to prevent unauthorised alterations to AI systems or models and ensure changes to AI (such as underlying code, configurations, or training data) do not result in unintended consequences such as service outages, operational disruptions, or degradation in model performance. This may be supported by version control systems that track changes not only to the AI model code, but also

⁴⁹ E.g., changes to training data, model or system architectures, key assumptions, scope of use, or possible impact of the change on downstream AI systems or models.

to associated data, parameters, hyperparameters, and other key artifacts, and support traceability, auditability, and the ability to roll back to previous versions.

- (c) **Dynamic Updates:** Implement enhanced controls for AI that are designed to be updated automatically, including strict justifications for enabling automatic updates, clear definitions of what can be updated automatically (such as allowing re-training or changes to hyperparameters but not changes to core architecture), enhanced data quality checks, and more stringent performance monitoring.

5.26. **An FI should develop and implement clear controls for the eventual retirement or decommissioning of AI when they are no longer needed or exceed risk tolerances.** These controls should consider dependencies, data retention policies, secure removal from the production environment, and appropriate stakeholder notifications.

6. AI Capability & Capacity

AI Risk Management Capabilities

- 6.1. **An FI should determine and ensure the necessary competence and proper conduct of personnel involved in developing, deploying and using AI.** This includes recruiting the necessary talent and providing appropriate training and support to equip employees with the skills, knowledge, and culture for the appropriate use of AI, as well as effective AI risk management. Adequate resources, including human, technological, and financial, should be allocated proportionate to the risk profile of the AI use case and their underlying systems or models.
- 6.2. **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.** This includes regularly reviewing programmes for building capabilities and capacity for effective AI risk management to ensure that they are up to date and incorporate training to address risks associated with newer AI technologies.

Technology Infrastructure for AI

- 6.3. An FI should ensure that its technology infrastructure is adequate for the AI use cases, systems or models it wants to implement. Like other information technology systems, AI systems and models used in such systems are also subject to technology risks, including system availability, resilience, safety and cybersecurity risks. To address these risks, the FI should ensure that the underlying hardware and software resources, such as graphics processing units, network infrastructure, system memory, secure data pipelines are sufficient to meet performance, scalability, and resilience needs for the AI use case, system or model. In determining the appropriate technology infrastructure for a specific AI use case, system or model, the FI should take into account relevant technology risk management guidelines and notices⁵⁰, as well as relevant industry frameworks⁵¹.

⁵⁰ <https://www.mas.gov.sg/regulation/guidelines/technology-risk-management-guidelines>

⁵¹ For example, technology infrastructure related areas covered in the NIST AI Risk Management Framework at <https://doi.org/10.6028/NIST.AI.100-1>.