asifma
Growing Asia's Markets

30 Jan 2026

To:
Alan Lim
Director & Head
Financial Infrastructure & Artificial Intelligence Office
Monetary Authority of Singapore

ASIFMA Response to MAS Consultation Paper on Proposed Guidelines on Artificial Intelligence Risk Management for Financial Institutions

Dear Mr Lim,

The Asia Securities Industry and Financial Markets Association ("**ASIFMA**")[1] would like to thank MAS for initiating this important consultation on the Guidelines for Artificial Intelligence ("**AI**") Risk Management (the "**Guidelines**") and welcomes the opportunity to provide feedback on behalf of our members. We strongly support MAS' objective of promoting responsible AI adoption while safeguarding financial stability, consumer protection and market integrity. A technology-neutral, risk-based, principles-based framework is essential to enable innovation without imposing unnecessary and/or undue compliance burdens.

General Comments:

We would strongly support alignment of the Guidelines with the OECD definition of an "AI system" and avoiding the interchangeable use of the terms "use case," "model," and "system." Clear terminology is critical for effective supervision and risk management. We also emphasise the high importance of distinguishing between high-impact and/or client-facing AI and low-risk and/or internal AI, based primarily on how the system is used and its potential impact. Higher-impact AI warrants a more stringent framework and oversight, whereas lower-risk applications do not require these and can be governed within existing frameworks.

We also wish to make overarching comments at the outset, highlighting the critical importance of:

- Allowing institutions to leverage existing global or regional AI governance and risk management structures, to avoid duplication and fragmented oversight;
- Providing flexibility in assessing AI risk materiality, recognising that risk arises from how AI is used, the degree of reliance, and potential impact—not from model complexity alone. We urge caution against using complexity as a primary driver of risk categorisation, as it is not a reliable

[1] ASIFMA is an independent, regional trade association with approximately 150 member firms comprising a diverse range of leading financial institutions from both the buy and sell side, including banks, asset managers, law firms and market infrastructure service providers. Together, we harness the shared interests of the financial industry to promote the development of liquid, deep and broad capital markets in Asia. ASIFMA advocates stable, innovative, and competitive Asian capital markets that are necessary to support the region's economic growth. We drive consensus, advocate solutions and effect change around key issues through the collective strength and clarity of one industry voice. Our many initiatives include consultations with regulators and exchanges, development of uniform industry standards, advocacy for enhanced markets through policy papers, and lowering the cost of doing business in the region. Through the GFMA alliance with SIFMA in the United States and AFME in Europe, ASIFMA also provides insights on global best practices and standards to benefit the region. More information about ASIFMA can be found at: www.asifma.org.

ASIFMA Headquarters
Unit 3603, Tower 2, Lippo Centre
89 Queensway
Admiralty, Hong Kong
Tel: +852 2531 6500

ASIFMA Singapore Office
12 Marina View, Asia Square Tower 2
Level 11, #11-01
Singapore 018961

nor future-proofed measure, and could impose unnecessary burdens and impede innovation. Oversight should instead be anchored in context, use, and impact;

- Alignment with and reliance on existing MAS frameworks, including the Technology Risk Management Guidelines ("**TRMG**"), to avoid parallel, duplicative processes and enable proportional, feasible implementation, especially where AI tools are embedded in non-AI specific services. Equally, new governance structures should only be created where absolutely necessary; and
- Focusing the Guidelines more on third party risk management where AI is procured directly from a supplier, than where it is used by a third party in the provision of a service (which would capture a very broad range of third-party relationships) – noting that in the latter case, further dialogue may be appropriate on the question of allocation of responsibilities.

We believe that ASIFMA's AI[2] and Generative AI[3] white papers will help to inform MAS' proposed approach. Both white papers emphasise leveraging existing regulations, applying proportionate oversight and avoiding blanket requirements that hinder efficient deployment. We support MAS' overall approach, particularly its emphasis on proportionality and life-cycle based risk management—consistent with global best practices.

In addition to our general comments, please also find our responses to the questions set out in the consultation paper in the Annex.

ASIFMA appreciates MAS' efforts to translate high-level principles into practical supervisory expectations that balance innovation with strong governance. While we broadly support the proposed framework, we offer targeted comments to ensure operational practicality and global consistency. We look forward to continued engagement with MAS to support a harmonised, resilient and trusted approach to responsible AI adoption across the financial sector.

We would be pleased to discuss our response in further detail. Should you wish to do so, please do not hesitate to contact me at rkapoor@asifma.org.

Sincerely,

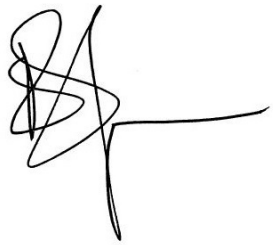

Rishi Kapoor
Executive Director
Head of Technology and Operations
Asia Securities Industry & Financial Markets Association

[2] https://www.asifma.org/wp-content/uploads/2021/06/enabling-an-efficient-regulatory-environment-for-ai-report_june-2021.pdf
[3] https://www.asifma.org/wp-content/uploads/2024/01/2024-asifma-gen-ai-paper-final-updated-18032024.pdf

ANNEX

Question 1: MAS seeks comments on the application of the Guidelines to all FIs in a proportionate manner, and the guidance on the proportionate application set out in paragraph 1.5 and the Annex of the Guidelines.

1.1 Broad Scope and Interpretation and Proportionality Challenges:

a. We believe that the Guidelines' scope is overly broad, potentially requiring full AI governance for low-impact internal tools (e.g. anomaly detection), which seems disproportionate. The framework should more explicitly and more consistently apply proportionate and risk-based principles. MAS should confirm that expectations align with the risk of the AI application.

b. We are concerned that while the Guidelines reference risk materiality, they do not consistently apply proportionality across lifecycle controls. Expectations should explicitly scale with risk arising from the use of AI, ensuring more intensive governance only applies to AI uses with significant impact, critical reliance or material consequences of failure.

c. We suggest that MAS aligns with international standards (e.g. OECD – see our General Comments) and adopts a principles-based, risk-proportionate approach that evaluates AI applications based on their overall risk materiality rather than fixed categories, similar to the approach taken by Hong Kong and other Asia-Pacific regulators. This includes distinguishing appropriately between higher-risk tools and lower-risk tools, with governance obligations scaling accordingly. Factors such as whether a tool is internal vs external or client-facing vs non-client-facing should inform the risk assessment but not automatically determine it. This would ensure targeted focus on material AI use cases that have high impact, while avoiding blanket requirements that hinder low-risk, assistive AI innovation.

d. Additionally, risk is determined by AI usage, not the AI model, hence AI models, systems and use cases should not be used interchangeably.

e. Applying the Guidelines to third party services that embed AI features would result in a very wide range of suppliers, given the broad and ever-increasing use of AI. In addition to capturing an overly broad scope, it is unclear how the full set of AI risk management expectations could meaningfully or proportionately be applied in all such cases. Most FIs would therefore suggest considering limiting the scope to AI that is procured by the FI, rather than extending it to AI used by suppliers as part of their operations or the delivery of the service.

1.2 Comparison, Harmonisation with Other Frameworks and Existing Guidelines

a. We suggest that the Guidelines should follow the precedent set in the MAS Outsourcing and TRMG and Notices, applying more rigorous obligations only to material or critical AI use cases that could affect regulated activities, disrupt services or significantly impact clients, while allowing non-material AI uses to be monitored with lighter requirements. AI lifecycle controls should also be explicitly integrated with existing MAS frameworks (e.g. TRM, Outsourcing, Cyber Hygiene), to ensure coherence and avoid duplicating high-level concepts, processes and obligations.

1.3 Clarification to Non-Integrated AI Use

a. There is uncertainty about the treatment of non-integrated AI use: while paragraph 3 of the Annex suggests only basic AI policies are needed, footnote 13 indicates that lifecycle

controls may still apply. We seek clarification from MAS on the intended scope for non-integrated AI under the Guidelines.

b. For FIs with no AI usage, or where AI is not an integrated part of business processes, we suggest that the expectations in paragraph 5 of the Annex be positioned as illustrative rather than prescriptive, and implemented in a manner commensurate with the risks identified. This would support a proportionate and practical approach across a diverse range of FIs.

1.4 Balanced View on AI

a. We believe the Guidelines would benefit from a more balanced perspective on AI. As currently written, they emphasise its risks, while offering only a brief mention of its benefits (e.g. in paragraph 1.8), which could unintentionally suggest that MAS discourages AI adoption. We recommend highlighting AI's benefits, showcasing successful implementations by Singapore-based FIs and reiterating MAS' overarching aim to drive AI adoption across the financial sector along with proper governance, as per the Guidelines.

Question 2: MAS seeks comments on the proposed scope of AI use cases, systems and models for the application of the Guidelines.

2.1 Clearer Definition

a. We note that footnote 33 of the Guidelines suggests that a high-risk use case should be regarded as using AI as an integrated part of business processes, however it does not explicitly exclude business-process-integrated AI with minimal risks (such as basic automation, non-generative uses, extractive tasks and structured data annotation). This ambiguity could lead to the inclusion of low-risk internal tools under the same governance framework, undermining the principle of proportionate application. Therefore, for the avoidance of doubt, members would recommend an explicit clarification that low-risk yet business-process-integrated AI is not considered high-risk.

b. Further, we note that the same footnote emphasises use cases deployed in critical business lines or regulated activities, which should be the anchor for defining scope and materiality. In this context, we recommend that MAS clarifies whether the use of AI in relation to the conduct of regulated activities should automatically be deemed high-risk AI.

2.2 Clarity on Exclusions and AI Examples:

a. Though we appreciate the intent behind the Annex's examples of what AI use cases are in scope and what are not, the list in paragraph 4 of the Annex cannot possibly be exhaustive. We suggest adopting clearer, principles-based criteria to define in-scope AI—ensuring that those examples serve only as illustrations—and that internal and/or low-impact tools are not subject to the same scrutiny as client-facing and/or decision-relevant AI. Members believe that the scope should be anchored to materiality and client or regulatory impact, with AI influencing client decisions or regulated activities clearly in scope, while internal or indirectly related applications (such as IT helpdesk bots or contract review tools referenced in section 4(b) of the Annex) are treated as lower-risk with minimal obligations, avoiding disproportionate baseline requirements.

b. More generally, members feel that the exclusions in the Guidelines in paragraph 3 of the Annex are not sufficiently clear - for example, why certain AI applications such as research

paper summarisation are excluded. We feel that the Guidelines should be risk-based, with criteria that may include internal vs external-facing generated content. If the tools can generate client-facing content, they should be regarded as being an integrated part of the business process.

2.3 Materiality and Integration:

a. We seek clarity regarding what constitutes "materially dependent" AI in paragraph 2(a) of the Annex, and how this threshold is defined, particularly in relation to business processes. However, we also caution against introducing definitions or thresholds which risk undermining FIs' ability to assess their material dependencies in a way that reflects their operating contexts. Therefore, principles-based guidance at a high level which respects existing risk management processes may be the most appropriate approach.
b. We also have concerns over the proposed material impact assessment in section 4.4. This should be based on existing frameworks and assessed based on actual impact to clients and services, not on the complexity of the technology type, whether AI or otherwise.
c. We seek clarity on what is deemed a "material AI risk" with examples or criteria (e.g. criticality to operations) that would warrant a dedicated AI committee or additional oversight layers.

2.4 Addressing Modern AI Systems:

a. We note that although MAS acknowledges AI agents, the Guidelines still appear to rely on traditional machine learning assumptions. Given the rapid evolution of LLMs and agentic systems, we suggest keeping the Guidelines risk-based, principles-driven, and outcomes-focused to ensure they remain relevant as technology advances. Prescriptive rules may quickly become obsolete.

2.5 Amendment to Existing Provisions

a. We suggest generalising certain risk descriptions, such as revising section 6, paragraph 1.8(a) to more broadly state that AI-related uncertainty in financial risk management may lead to poor assessments and financial losses. A more general formulation would better reflect the inherent complexity of AI risks and keep the framework adaptable as technologies evolve.
b. We also suggest refining the terminology in section 6, paragraph 2.3(a) by replacing "relevant" with "material", as this provides a clearer, more focused threshold and avoids an overly broad interpretation.

Question 3: MAS seeks comments on the proposed responsibilities of the Board and senior management in overseeing AI risk management.

3.1 Granularity of Board Responsibilities:

a. We believe that the responsibilities of the Board outlined in section 2 of the Guidelines may be overly granular. We suggest clarifying that the Board should set the overall tone, strategy, and risk appetite frameworks, while detailed implementation, operational execution and day-to-day oversight within the approved risk appetite are delegated to Board committees and senior management. This better reflects the typical division of responsibilities.

3.2 Training and Inclusion of AI Expertise to Higher Management

a. We support the benefits of training and/or having appropriate AI knowledge at the board/senior management level, to ensure that the board can fulfil the responsibilities set out in our remarks in section 3.1(a) of this response (similar to section 3.1.2/3.1.3 of the TRMG).

Question 4: MAS seeks comments on the proposal for FIs to establish a dedicated cross-functional committee to oversee AI risk if the overall AI risk exposure of an FI is deemed material; and how such overall AI risk exposure should be assessed at the organisational level.

4.1 Proposed Committee Structure

a. Members firmly believe that the requirement to establish a new, dedicated committee to provide holistic, cross-functional oversight of AI risks across privacy, security, cloud, and emerging technology domains is unnecessary and should not be mandated, as most FIs already have established, effective cross-functional governance bodies (e.g. Risk and Cybersecurity Committees) and/or senior risk officers capable of overseeing AI risks. Since the intent of MAS appears to be ensuring cross-functional oversight of AI risk rather than mandating the formation of a new, standalone committee, FIs should retain the flexibility to leverage existing structures rather than be forced to create new ones, ensuring proportionality, keeping AI governance manageable (particularly for smaller firms) and avoiding unnecessary burden.

4.2 Regional Flexibility

a. We believe that global or regional committees would be entirely appropriate to leverage for the oversight purposes above, as many parent entities develop AI tools and their related risk management and governance frameworks centrally for cross-jurisdictional use. Leveraging these structures allows members to utilise effective, experienced, existing mechanisms, maintain visibility, enable effective escalation and ensure consistent oversight without duplicating governance locally.
b. We suggest that MAS confirms that FIs may rely on parent-level or global AI risk management frameworks and governance structures, as contemplated in paragraph 3.1 of the consultation paper. This includes confirming that global boards, senior committees, regional governance forums and other governing bodies may fulfil MAS's expectations, provided the FI applies the Guidelines proportionately and demonstrates appropriate local adaptation, especially for cross-jurisdictional AI use. Broader terminology (e.g. "governing bodies," "management mechanisms" and "forums") and explicit recognition of global governance frameworks would help to clarify this point.
c. We feel this flexibility should extend to roles as well, including allowing the designated senior risk officer to be a regional appointee rather than a Singapore-based employee, where this aligns with global or regional AI governance structures.
d. We further suggest clearer articulation of the division of responsibilities between the Board and senior management, particularly when global or regional structures are leveraged, to ensure alignment with MAS' expectations.

Question 5: MAS seeks comments on the proposal for FIs to establish clear definitions, criteria and processes, supported by robust systems, to facilitate the consistent identification of AI usage across all relevant business and functional areas.

5.1 AI Identification by Control Function:

a. We suggest that for AI Identification (paragraph 3.2), there should be mention of the control function if the intent is to be aligned with the assignment of roles and responsibilities for AI identification (paragraph 3.3).
b. We seek clarification on what should be the designation of the control function responsible for AI identification systems and processes (paragraph 3.3), in relation to the involvement of other existing, relevant control functions that already support the identification of AI.
c. We seek clarification on the setting up of attestation processes (paragraph 3.3) as to whether this is in relation to the roles and responsibilities, application of AI identification or otherwise.
d. We seek clarification on the expected level of automation for the AI identification process—specifically, whether it should rely on automated discovery tools, manual attestations or a risk-based combination of both.

Question 6: MAS seeks comments on the proposal for FIs to establish and maintain an accurate and up to-date inventory of all AI usage.

6.1 Administrative Burden:

a. We strongly suggest adopting a proportionate, risk-based approach to the inventory requirements in the sections related to AI use case tracking, as maintaining a comprehensive register with full attributes for every AI use case, model, system and/or third party feature could create immense operational burden, substantially change the cost-benefit analysis of potential new adoptions and act as a disincentive to future take-up, given the expanding use of AI. The focus should be on material AI uses that could meaningfully affect customers or operations, similar to the materiality-based registers in the MAS TRM and Outsourcing Notices, where 'higher-risk' use cases include the full set of required attributes, while lower-risk cases require only a minimal subset.
b. We also believe it should be acceptable for the inventory/inventories to be maintained at the global or regional level, as long as the requirements in the Guidelines are fulfilled, as many parent entities already centralise AI development and maintain a single, global inventory with regional input. This allows access to cross-functional expertise, efficient prioritisation and controlled roll-out, while avoiding unnecessary duplication in Singapore.

6.2 Pragmatic Challenges with Inventory Management:

a. We wish to highlight several pragmatic challenges:
    1. The Guidelines mention exclusions (e.g. simple research summarisation), however it is unclear whether these exclusions also apply to the inventory requirement.
    2. There is ambiguity regarding whether multiple AI agents operating on a single underlying AI system would need to be inventoried separately, which could lead to an unmanageable inventory.
b. We suggest that MAS provides illustrative examples of the scope of the inventory (e.g. internal AI use cases, third party usage of AI in service provisions, etc.). This could include provisions in the inventory such as how to treat agents versus underlying systems and how to align exclusions with inventory obligations.

c. We seek guidance on what counts as Singapore-relevant AI usage, such as AI supporting Singapore-booked clients, AI used by Singapore-based staff or AI embedded in global platforms with limited relevance to Singapore. Clear definitions would help drive consistent implementation and avoid over-capture.

6.3 Guidance and Feasibility of AI Inventory

a. We would like to inquire if MAS is intending to produce an optional template/decision tree/guidance to guide FIs, particularly for borderline or emerging use cases, to support consistency across the industry. While noting that members have mixed levels of support for such an initiative lest it becomes a prescriptive tool, they would certainly request timely engagement with industry at the design stage and well ahead of publication, if MAS has such an intent to produce one.
b. We suggest that FIs should implement a systematic process for AI identification, assessment and inventory, and allow the Three Lines of Defence model to operate as it does for other risk management processes.

Question 7: MAS seeks comments on the proposed risk dimensions of impact, complexity and reliance that should be captured by FIs in AI risk materiality assessments, and whether there are any other risk dimensions that should be included.

7.1 Unreliability of Complexity:

a. Most members believe that complexity should not be treated as a primary risk-tiering dimension, as complexity does not scale well as AI technologies evolve, nor does it reliably correlate with real-world risk outcomes. A simple model that materially affects clients may be high-risk, while a complex internal tool may be low-risk. Additionally, classifying all AI agents as inherently "complex" weakens the usefulness of complexity in risk assessment. As stated previously, risk should be driven primarily by the impact of the use case and the degree of reliance, not the technical sophistication of the model or system. System or model characteristics and novelty should only be treated as secondary modifying factors, not standalone drivers.
b. We also note global regulatory trends which are de-emphasising complexity as a core risk dimension.
c. We observe that paragraphs 3.8 – 3.11 describe criteria for assessing materiality rather than defining materiality itself. It is unclear whether this is intentional.

7.2 Focus on the Use of AI System

a. As stated earlier, AI risk materiality arises from how the system is used, the level of reliance placed on it and its potential impact on customers, markets, and critical operations—rather than from the model or system in isolation. The Guidelines should therefore allow firms flexibility to assess AI risk materiality using existing risk management frameworks, explicitly recognising that materiality depends on use, reliance, and consequences of failure. This aligns with MAS's Outsourcing Guidelines, which focus on assessing the materiality of the third party relationship (including AI models), and avoids introducing a separate, duplicative materiality assessment that would add complexity without benefit.

Question 8: MAS seeks comments on the proposed standards, processes and controls that should be applied across the entire AI life cycle, and the key areas that FIs should assess for relevance to the AI model, system or use case, and apply in a proportionate manner.

8.1 General AI Life Cycle Controls:

a. Per Principle 1 of the ASIFMA Generative AI white paper, we strongly suggest leveraging existing data management, procurement and outsourcing frameworks, including embedding third party AI risk management within current vendor governance arrangements.
b. For clarity, we would support a consistent application of a risk-based approach through all the life cycle controls, including data management expectations, selection, reproducibility and auditability.

8.2 Data Management, Evaluation and Testing:

a. We welcome reliance on existing data governance standards, but suggest clearer proportionality in data expectations, particularly recognising limited transparency over training data for third party models. The phrase "full range of real-world conditions" would benefit from a tighter definition to avoid ambiguity.
b. We also seek clarification on data usage permissions, including whether explicit consent from data subjects is required for training third party models, and how these expectations align with existing MAS guidance such as the Data Governance & Management Practices Information Paper 2024.
c. We seek clarification on footnote 14 in the Guidelines, related to what minimum controls are expected before a pilot/partial deployment can occur (e.g. baseline testing, threat assessment, data controls, human oversight), and what triggers the requirement to complete full pre-deployment independent validation.

8.3 Transparency and Explainability

a. Practices related to transparency and explainability continue to evolve. MAS should provide flexibility to reflect evolving industry practices, and that output-based monitoring controls (e.g. back-testing, anomaly detection) may be appropriate in certain instances.
b. We propose that MAS offers clearer guidance on the disclosures required when AI is used in customer-facing processes, including what information must be communicated, whether model features or data attributes need to be referenced, and what form an appropriate customer redress channel should take.

8.4 Fairness

a. We also request clearer definition of "protected attributes" in the Singapore context, guidance on handling proxy variables, and more concrete direction on fairness testing scope and monitoring expectations.

8.5 Third Party AI Management – Scope of Application

a. Most FIs suggest that the scope of application is limited to third-party procured AI. The references to third party AI under footnotes 10, and 21 and paragraph 4.11, being the current broad definition—which covers procured AI, supplier-used AI and services where AI is later introduced—suggests that a wide range of third party AI scenarios are intended to be in scope. As most third party services are expected to incorporate AI tooling into non-AI

specific services, this could result in a massive expansion of the scope of applicability to third party services that should not be – and likely cannot be from a practical perspective – subjected to the type of governance and oversight set out in the draft Guidelines.

b. Whilst we recognise the rationale for capturing third party AI, some of the specific governance and lifecycle controls – such as the risk materiality assessment, model-level oversight and certain technical controls – do not readily map to all forms of supplier use of AI, and it is not clear that these requirements were designed with the full range of third party scenarios in mind.

c. If MAS does intend to include supplier use of AI broadly, clarifying how these expectations are intended to apply in the context of supplier usage of AI, as distinct from procured AI solutions, would support a more consistent and proportionate application. In particular, expectations should reflect the FI's degree of control and visibility over the AI capability. Firms should not be expected to meet lifecycle obligations where the AI operates entirely within the supplier's environment and is not procured, configured or managed directly by the FI.

d. While FIs should retain full responsibility for internally developed models, FIs also take the view that third party providers should remain accountable for elements under their control, and propose that a dialogue should be started to discuss how responsibilities should be assigned between deployers and external providers. MAS may wish to clarify how responsibilities for risk management and controls can be appropriately allocated in practice, based on the AI deployment model (e.g. on-premises, PaaS, SaaS) and consistent with existing cloud best practice guidance in the paper "Cloud Best Practices for Singapore Financial Institutions". AI risk management requirements should reside with the entities best placed to address them, supported by clearer expectations on supplier assurances (e.g. attestations and certifications). Similar approaches were explored in MAS' Veritas Initiative and the EU AI Act.

8.6 Third Party AI Management – Risk Management

a. We suggest aligning the Guidelines with the MAS TRMG and Outsourcing Notices to ensure integration with existing frameworks. Third party risk management (**TPRM**) frameworks are designed to assess risk at the arrangement level and adapt to underlying technology risks (data, security, resilience), taking into account the materiality of the arrangement.

b. Materiality in this context focuses on the potential impact of failure (as set out in Annex 1 of the Guidelines on Outsourcing), which also considers its connectivity to critical systems (as referenced under the TRM Notice), among other factors. The AI-specific materiality test—which focuses on model/system characteristics—creates a parallel and duplicative framework. This creates complexity for FIs' TPRM frameworks and assessments, and raises questions around how these requirements are intended to integrate with established TPRM frameworks and technology-agnostic regulatory requirements.

8.7 Third Party AI Management – Notification & Controls

a. FIs are concerned about frequent, unnotified changes to third party AI products, which complicate compliance with update notification and impact assessment requirements as outlined in section 6, paragraph 4.11. We strongly suggest that notification requirements instead focus on directly contracted general-purpose AI providers (e.g. foundational models such as ChatGPT), rather than all vendors that may incidentally introduce AI features.

b. More broadly, most FIs are concerned that the current drafting may inadvertently capture third party services where AI use is incidental or supplementary, and suggest limiting the

scope of paragraph 4.11 to services explicitly marketed as AI tools for advisory or customer-facing use or those which have high impact, to avoid excessive compliance and administrative burden on FIs.

c. We seek clarification in paragraph 4.11 on what "development" controls mean in a third party context: whether they refer to the provider's AI development processes and standards, or to broader ongoing governance and oversight of third-party AI systems.
d. We recommend that MAS provides clarity on the interpretation of the terms "introduced" and "incorporating" AI, including whether this extends to background use of AI (e.g. copilots). Additionally, we request MAS to define a threshold for notification, such as when the AI poses risks to the fundamental rights of natural persons.
e. Lastly, we seek clarification on "performing compensatory testing to address informational gaps arising from inadequate disclosures by third party providers" in section 4.11 of the Guidelines. Particularly, we ask how 'inadequate' is defined and how it is applied to commercial models (and/or open-source models, or specific vendor solutions that are SaaS).

8.8 AI Selection

a. We believe that in paragraph 4.12, documentation should not be mandated for low-risk applications, as it does not necessarily lead to better assessments and can impose overly burdensome processes.
b. We propose that under paragraph 4.13, AI algorithms piloted on a small scale should be permitted to move to deployment before the evaluation and testing phases are fully concluded. We suggest that low-risk use cases may not require the complete set of evaluations, testing or overfitting techniques.

8.9 Technology and Cybersecurity Risks

a. We request MAS to clarify the "kill switches" expectations, as they may not be feasible for some architectures. For 'high-risk' AI, clarity would be beneficial for members in terms of when MAS expects a true kill switch (immediate disablement) vs other controls (rate limiting, fallback routing, feature flags, scoped disablement) to be applicable.

8.10 Reproducibility and Monitoring

a. We agree with MAS' focus on explainability, but caution that the current reference to reproducibility of "results" is not realistic for traditional AI/ML models, and will present a greater challenge for emerging models such as reinforcement learning or continuously updated models. We suggest shifting the focus from replication to the processes, templates, and auditability of methodologies, and broadening the language to accommodate emergent errors and behavioural drift in generative models.
b. As noted above, we suggest that the Guidelines include a consistent reference to a risk-based approach through all the life cycle controls, including reproducibility and auditability.
c. In addition, we also suggest that paragraph 4.7 referencing specific hardware (e.g. GPU, TPU) be made more technology-neutral and focused on risk outcomes.

8.11 Reviews and Monitoring

a. We seek further clarity regarding paragraph 4.23(a) on monitoring measures: for instance, illustrative examples of metrics necessary for effective implementation and operation by the technical team.

b. While we acknowledge that MAS allows for a risk-based and proportionate approach in both paragraphs 4.16 and 4.22, the current wording in both sections may be interpreted as strict requirements. We recommend that MAS explicitly incorporates proportionality into these sections to enhance clarity and flexibility. Specifically, for paragraph 4.22, we suggest modifying the text to state: "Such reviews can take into consideration and cover areas such as technical implementation, system and network security, system resilience, recoverability and operational readiness of AI for deployment," while removing the subsequent provisions under the paragraph.

c. Lastly, we seek clarification on the term "independent parties" as mentioned in paragraph 4.24, requesting confirmation that the assessment of independent parties can be based on Operational Performance Management (**OPM**) results and that periodic revalidation may not be required. Furthermore, we recommend that MAS explicitly states that "independent parties" may include internal teams that operate independently from the AI execution teams. This clarification would improve understanding and facilitate proper adherence to the Guidelines.

8.12 Stress Testing

a. We seek clarification on "stress testing" (section 1.6 of the Guidelines) in the context of whether it is capacity-related or technical in nature, or determined by the ability to understand how the AI system functionally behaves under pressure (e.g. asymptotic prompts).

Question 9: MAS seeks comments on any aspects of the Guidelines that have not been covered in earlier questions, as well as aspects of AI risk management that have not been covered in the proposed Guidelines.

9.1 Amendment to Existing Provision

a. We request further guidance on multi-agent or complex AI systems, with greater direction on designing, measuring and monitoring the risks and controls for each agent. Additional guidance to cover the sensitivity and amplification of risks through the chain of agents and the overall outcomes from the system would be beneficial.

9.2 Technical Infrastructure Guidance:

a. We believe that paragraph 5.3 of the Guidelines is overly granular and too prescriptive in specifying technical infrastructure (e.g. processor-type references). Given the rapid evolution of AI technologies, MAS may wish to adopt more principles-based statements—such as requiring infrastructure to be "resilient and scalable"—rather than detailing chip types or architectures. We further recommend aligning the section 5.3 requirements with the existing TRMG and related notices to ensure integration with established frameworks and to support proportionate, risk-based implementation.

Question 10: MAS seeks comments on the proposed transition period of 12 months.

10.1 Transition Period and Resource Utilisation for AI Guidelines

a. Members do not object to the proposed 12-month transition period, provided that MAS offers checklists, toolkits and templates for self-assessment during this time. We suggest

that MAS leverages existing checklists and templates from the TRM framework by incorporating AI-specific considerations, such as model transparency and third party AI risk, to facilitate self-checking for firms and confirm compliance readiness before full implementation.

b. We also recommend that MAS offers more detailed guidance and tailored templates for smaller FIs, particularly for AI inventory, materiality assessment and control mapping, to minimise unnecessary compliance burden.

c. We recommend that MAS also considers an extended or phased implementation timeline, given the significant breadth of the Guidelines and the need to align with multiple global regulatory regimes. This could include allowing additional time for more complex areas or adopting a phased approach for key controls such as AI inventory, lifecycle management and third-party governance.