Annex B

Submissions from respondents to the consultation paper on Guidelines on Artificial Intelligence Risk Management

Note: The table below only includes submissions for which respondents did not request confidentiality.

S/N	Respondent	Responses from Respondent
1	/TMP Labs Pte. Ltd. (Singapore)	**Question 1:** We support MAS' emphasis on proportionality. In practice, proportional application depends on an institution's ability to maintain visibility over AI usage across diverse contexts, from low-risk internal productivity tools to high-impact decisioning systems. Lifecycle- and context-aware application is more effective than static classification. **Question 2:** We agree with a broad scope covering both internally developed and externally sourced AI systems. Increasingly, risk arises from composite and agentic systems where multiple models and tools interact, rather than from single models in isolation. Recognition of such systems is important for practical risk management. **Question 3:** Clear accountability at the Board and senior management level is appropriate. Operationally, this accountability is most effective when supported by reliable, timely information on AI usage and change, rather than reliance on periodic reporting alone. **Question 4:** Cross-functional committees are valuable, provided they are supported by a shared and up-to-date view of AI inventory, ownership, and usage. Information coherence across business, technology, risk, and compliance functions is critical to effective governance. **Question 5:** We support lifecycle-based risk management. Early-stage discovery and onboarding of AI use cases is particularly important, as AI adoption often precedes formal approval processes. Ongoing change awareness is essential given the evolving nature of AI systems. **Question 6:** We view an accurate and current inventory of AI usage as foundational. Practical challenges include shadow AI adoption and AI capabilities embedded within third-party platforms. Maintaining inventory should be treated as a continuous operational process.

		Question 7: Risk scoping is most effective when grounded in operational context. Clear separation between evidence collection, risk assessment by qualified personnel, and regulatory determination helps preserve accountability and avoids ambiguity in governance. **Question 8:** We support proportional and context-driven application of standards and controls. Clear boundaries between documentation, assessment, and assurance outcomes reduce the risk of over-reliance on tooling and preserve professional judgment. **Question 9:** Agentic AI systems and autonomous workflows introduce emerging operational risks. Guidance that emphasises traceability, change awareness, and human oversight—without prescriptive technical requirements—would be helpful as these systems evolve. **Question 10:** A transitional period is appropriate. Institutions may require time to establish foundational visibility, inventory, and accountability structures before advancing to more mature AI risk management practices. We appreciate MAS' consultative approach and support the principles-based, proportionate framework proposed. We thank MAS for the opportunity to provide comments.
2	Access Partnership	**Question 1:** Access Partnership supports MAS' intention to apply the Guidelines to all FIs with proportionality. However, we would recommend: - Clearer thresholds or decision criteria for when a Financial Institution (FI) is considered to "use AI as an integrated part of its business processes" (beyond the examples in the Annex). - A simple self-assessment template or flowchart (building on Figure 2 in the Annex) that supervisors and FIs can both use to achieve consistent classification. - Explicit confirmation that FIs may rely on group-wide frameworks (subject to local gap-filling) to avoid duplication for branches and subsidiaries. The proportionality concept is good, but the current drafting still leaves grey areas (e.g., what counts as "integrated", where to draw the line between pilots vs production, and how far down the supply chain to go for embedded AI in third-party systems). Without sharper criteria, proportionality will be applied differently across firms and inspections. Many FIs already have group AI/model governance, so allowing reliance avoids fragmentation while still meeting MAS expectations. **Question 2:**

		Access Partnership supports the broad, technology-neutral definition of AI (models, systems, and use cases that learn or infer from inputs) and exclusion of purely rule-based calculators. However, we recommend clarifying the following for the FIs: - How do we consider mixed systems that combine rule-based logic with minor statistical components, i.e., when do they "tip over" into AI? - RPA / workflow tools with embedded vendor models but no direct model control by the FI. - Simple statistical models (e.g., linear regressions, scorecards) used in traditional risk management – whether MAS intends these to be in scope if they are not adaptive post-deployment. - While "Marketing staff using AI image generators to assist in designing marketing materials" is excluded, this could lead to misleading claims for financial services. While the breadth is generally sensible and appropriate for futureproofing, some borderline technologies (advanced rule-based systems, rule-plus-simple-stats, RPA with ML components) will require clarification to avoid inconsistent classification and "over-capture". Consistent classification is vital for meaningful inventories and risk assessments. **Question 3:** Access Partnership agrees that Board and senior management should oversee AI risk, set risk appetite, and ensure AI use does not undermine other supervisory expectations (e.g., fair dealing). However, Boards may not yet have sufficient AI literacy to challenge management meaningfully. At the same time, many requirements (inventories, model validation, data governance, TRM, outsourcing due diligence, etc.) overlap heavily with existing requirements for models and technology. Therefore, we recommend MAS to: - Clarify that Boards can discharge their oversight via existing Board-level committees (e.g., risk, technology) rather than creating new ones in all circumstances. - Highlight the importance of Board-level AI literacy and training, and acceptance that Boards may rely on specialised internal functions and external experts for technical depth. - Clarify how existing risk appetite statements (e.g., model risk, operational risk) can be extended to include AI, to avoid building a separate AI-only risk appetite. Without explicit mapping, FIs risk building parallel AI governance structures instead of extending existing model risk, tech risk, and third-party risk frameworks. This way, FIs avoid governance proliferation while still ensuring robust oversight. This also reflects practical limitations in Board expertise and ensures realism while maintaining accountability. In addition, MAS may wish to add the responsibility of ensuring appropriate training and capacity building within the organisation to establish baseline AI literacy for all staff, and advanced/specialised AI expertise where needed. **Question 4:** Access Partnership supports the concept of a dedicated cross-functional oversight mechanism where overall AI risk exposure is material. We would recommend MAS to include:

		- More concrete guidance on how to assess "overall AI risk exposure" at the organisational level (e.g., number of high-risk AI use cases, use in critical business lines, reliance on external GenAI for key processes, client-facing vs internal). - Clarification that existing committees (e.g., Model Risk Committee, Operational Risk Committee, Technology Steering Committee) can jointly meet these expectations if their mandates explicitly incorporate AI. - For foreign bank branches and subsidiaries, recognition that a global or regional AI committee supported by local management may suffice, subject to local risk considerations. Without clearer triggers, FIs may either over-invest (creating committees for marginal use) or under-invest (arguing that exposure is not "material"). Building on existing structures reduces cost and complexity and leverages current expertise. **Question 5:** Access Partnership strongly supports the need for systematic identification of AI usage with clear definitions, criteria, and a central control function. To aid this further, we suggest MAS to: - Provide illustrative examples of identification criteria (e.g., self-learning, parameter updates post-deployment, probabilistic outputs, natural language generation). - Encourage a combination of top-down (policy-driven) and bottom-up (attestations, tooling such as API / SaaS usage scans) approaches to surface "shadow AI". - Clarify expectations on frequency of attestations and reviews, especially as new tools are rapidly adopted by staff. AI Identification is the foundation for everything else, and weak identification leads to blind spots in AI Risk Materiality Assessment and AI Inventory. Many AI tools enter organisations informally (free trials, productivity tools); a practical detection and attestation regime is needed. **Question 6:** Access Partnership agrees that an accurate, up-to-date AI inventory is essential, and supports MAS' suggested key attributes (purpose, scope, data, lifecycle status, roles, documentation links, etc.). We would, however, recommend clarifications on: - Granularity: whether inventory entries should be at the use case level, model level, or both (e.g., same model used in multiple use cases across jurisdictions). - Interaction with existing model inventories: encourage MAS to explicitly allow FIs to extend existing model risk inventories rather than maintain separate AI-only inventories. - Third-party / SaaS models: offer guidance on minimum metadata required when underlying model details are proprietary (e.g., category, provider, licensing arrangement, data flows, assurance artefacts) rather than detailed technical fields that may be unavailable.

		Without such granularity, AI inventories risk becoming unwieldy or inconsistent across entities. At the same time, leveraging existing inventories reduces duplication and implementation cost. **Question 7:** Access Partnership strongly supports the proposal to calibrate controls based on risk materiality, with the three dimensions of impact, complexity, and reliance as a minimum. We would also recommend: - Explicitly treating data sensitivity/privacy and third-party dependence/concentration as either separate dimensions or clearly embedded within "impact". - Clarifying that "complexity" should also reflect maturity and familiarity of the technology – i.e., a well-understood classical ML model may have lower complexity than a novel AI agent, even if parameter counts are similar. - Providing high-level examples of how to rate risk (e.g., low/medium/high) across the three dimensions for typical use cases (copilots vs credit decisioning vs AML monitoring). Data sensitivity and third-party dependence would be the key drivers of real-world risk and deserve explicit consideration. Guidance on rating avoids "risk inflation" or "risk under-statement" and facilitates comparability across FIs. **Question 8:** Overall, Access Partnership believes that the lifecycle controls are comprehensive and aligned with global best practice. However, we have some targeted comments: 1. Data management o We support strong expectations on fit-for-purpose, representativeness, quality, security, and privacy for data used across the AI lifecycle. We suggest: - Clarifying how proportionality applies to data governance for low-risk, internal-only use cases (e.g., internal productivity tools). - Recognising that for some vendor-provided models, FIs can only partly verify representativeness or training data quality and may need to rely instead on outcome testing and contractual assurances. 2. Transparency and explainability o We agree that explainability should scale with risk. AI used for credit, underwriting, or high-impact advisory decisions should be more explainable than internal copilots. We suggest MAS explicitly acknowledge: - That for highly complex models, post-hoc explainability methods and high-level explanations to customers may be acceptable. - That disclosure to customers about the use of AI can be layered (high-level for all, detailed on request). 3. Fairness o We support explicit fairness assessments where outputs affect customer access to products or pricing. We recommend: - High-level guidance or examples of fairness metrics and protected attributes relevant in the Singapore context (to avoid inconsistent or purely box-ticking approaches). 4. Human oversight

		o We agree with the emphasis on defined roles, competence, and documentation for human oversight. We suggest MAS explicitly highlight: - Risks of automation bias and over-reliance, especially for Generative AI and AI agents. - The need for training of front-line users to understand limitations and when to override AI. 5. Third-party AI management o We support strong expectations on diligence, transparency, supply-chain risk, and contingency planning. We recommend: - Harmonising expectations with existing outsourcing and third-party risk frameworks and perhaps providing a mapping table. - Acknowledging practical constraints where detailed training data and model architecture information cannot be disclosed, and emphasising outcome-based testing, independent certifications, and contractual safeguards as acceptable alternatives. 6. Selection, evaluation & testing; technology & cybersecurity; reproducibility & auditability; pre- and post-deployment controls o We support these as they are broadly aligned with MAS' existing TRM guidelines and model risk management paper. We suggest MAS: - Provide an example of a "minimum test pack" for low-risk applications vs a more comprehensive set for high-risk ones (e.g., adversarial testing, stress testing for GenAI and AI agents). - Clarify expectations for models that are continuously learning (online learning) vs periodically retrained, especially for change management and monitoring. - Encourage stepwise deployment (pilots, phased rollout) with clearly defined success criteria and limitations, as already alluded to for partial deployments. - Lifecycle controls are expected to be the bulk of the implementation cost. Therefore, clear examples of proportionality will prevent over-engineering low-risk use cases while still protecting customers and the system. **Question 9:** Access Partnership proposes the following aspects for consideration. 1. Customer redress and incident reporting o MAS could explicitly link AI incidents (e.g., systematic mispricing, discriminatory decisions, mis-selling via AI chatbots) to existing customer complaints and incident reporting frameworks, including expectations around remediation and ex-post root cause analysis. 2. Interaction with employment / HR use of AI o Many FIs are starting to use AI for recruitment, promotion, and performance management. These have fairness, ethics, and conduct implications, but are not explicitly mentioned. MAS may wish to clarify whether such internal HR use cases fall within the scope, especially when they potentially affect competence and conduct standards. 3. Cross-border and group-level coordination o Given that many AI services are hosted cross-border, MAS could emphasise how cross-border data transfer and extraterritorial AI regulations should be considered within the AIRG (e.g., EU AI Act, foreign sectoral rules), while avoiding duplication. 4. Alignment with industry frameworks (MindForge, NIST, ISO) o The paper already references Project MindForge and the forthcoming AI Risk Management Handbook. MAS could explicitly state that using such frameworks (e.g.,

		MindForge, NIST AI RMF, ISO/IEC standards) will be viewed as good practice when appropriately implemented. **Question 10:** Access Partnership supports the principle of a transition period and recognises that many FIs already have elements of AI governance in place. However, a 12-month transition period may be tight for large and/or AI-mature FIs that need to implement full AI inventories, risk-materiality methodologies, lifecycle controls, and group-level coordination, especially when group frameworks must be adapted to local expectations. Therefore, we recommend: o Allowing FIs to prioritise within the 12 months: - Phase 1 (first 6–12 months): AI identification, high-level inventory, risk materiality framework design, governance, and Board oversight. - Phase 2 (beyond 12 months, or on a "best efforts" basis): full lifecycle controls for lower-risk use cases, detailed integration with group frameworks, and broader capability-building. o Clarifying supervisory expectations that during the transition, FIs should at a minimum focus on high-risk AI use cases (those affecting customer outcomes, financial stability, or core risk management) even if lower-risk use cases are addressed later. A strict "all-or-nothing" requirement within 12 months could be challenging, particularly for FIs with complex global operations and numerous use cases. A phased, risk-based implementation is more consistent with MAS's proportionality principle.
3	AI Ling Advisory Limited	**Question 1:** **Question 2:** **Question 3:** **Question 4:** **Question 5:** **Question 6:** We agree with MAS that firms should maintain an accurate, current AI inventory linked to governance and risk management, with clear policies and responsibilities. We also note MAS's minimum attributes (purpose, approved scope, model type, data, dependencies, lifecycle status, risk materiality rating, validation status, roles, links) and the call to review the schema for newer AI and third-party attributes. Moreover, MAS highlights that unapproved usage can push the firm beyond risk appetite—hence the importance of secured environments and DLP. Our recommendation (aligned to EU AI Act risk tiering/minimal-risk practice) Two-tier inventory, calibrated by materiality. Tier A (Medium/High materiality): capture all MAS minimum attributes and apply independent validation proportionate to risk. Tier B (Minimal-risk embedded features): allow category-level registration (e.g., "productivity-suite AI features") where usage is confined to pre-approved corporate boundaries and mandatory guardrails are enforced (DLP, logging). Define escalation

		criteria (customer-facing, sensitive data, automated decisions, increased reliance) that promote Tier B → Tier A. This mirrors the EU AI Act's "minimal-risk" notion for low-risk features while preserving traceability. Recognise "secure AI sandboxes / walled gardens". Treat sanctioned internal AI environments (model allow-lists, safe-prompt filters, LLM-aware DLP/gateway, egress controls, central logging) as an accepted control to reduce unapproved usage, operating alongside the inventory. Outcome-based supervision for Shadow AI. Emphasise outcome metrics—e.g., sensitive-data DLP blocks to AI domains, policy-violation prompt detections, time-to-block SLAs—rather than inventory completeness alone; retain Tier-B→Tier-A escalation as risks rise. **Question 7:** **Question 8:** We support MAS's requirement for formal independent validation before deployment for AI assessed as high risk materiality, covering conceptual soundness, data suitability and quality, implementation integrity, testing results, explainability/fairness, and assumptions/limitations. We also welcome MAS's expectation that where third-party transparency is limited, firms conduct compensatory testing, maintain update/change notifications with providers, and manage legal clauses (performance, data protection, right to audit/notification). MAS's technology controls (secured environments, DLP) and the call for contingency planning further support a resilience-oriented approach. Our recommendation (aligned to EU AI Act provider/deployer roles) Provider attestations + FI "last-mile" controls. Where base models are consumed via API and provider transparency is limited, MAS should expressly allow reliance on provider-level evidence for the base model (e.g., safety/robustness evaluation summaries, change logs), while holding FIs accountable for application-layer controls (prompt/RAG design, data boundaries, output guardrails, adversarial testing). Independent validation then applies proportionately to the FI's implementation. This mirrors the EU AI Act's graduated, role-based obligations (provider versus deployer) and specific duties for foundation models. Outcome over anatomy. Assess compliance through resilience outcomes under FI control: drift monitoring with thresholds, kill-switch/override for high-risk use, and tested contingency/exit including provider switch readiness; codify material-change re-test triggers and provider notice SLAs. Minimum GenAI test-pack (baseline). For GenAI, name a minimum set: hallucination rate, toxic/biased content, data leakage/disclosure, and adversarial/jailbreak robustness—with recorded thresholds and mitigations. **Question 9:** **Question 10:**
4	Alternative Investment Management Association	**Question 1:** While we broadly agree that the proposed AIRG should be applicable to FIs in a proportionate manner, we would like to clarify whether "integration" is the appropriate trigger for requiring an FI to implement more rigorous policies to mitigate risks associated with the use of AI.

		The illustrative examples set out in paragraph 4 of the Annex to the proposed AIRG suggest that enhanced AI risk management measures are justified either because (1) the AI tool is widely adopted within the FI, or (2) the failure of the AI tool would result in widespread disruption. For example, tools with a high degree of “integration” may ultimately pose a very low risk of disruption to an FI's operations, whereas tools with a low degree of "integration" may conversely pose a very high risk of disruption. We propose for the MAS to consider adopting (2) as the primary test for requiring an FI to implement more rigorous policies, because it focuses the enquiry on the likely impact on the FI or its customers. In particular, AIMA suggests that the relevant threshold should be whether failure or malfunction of an AI tool would result in material disruption to the FI’s core business activities, customer outcomes, or regulatory obligations, rather than whether the tool is integrated across multiple workflows. This would be an approach that is similar to the test for material outsourcing arrangements under the MAS Guidelines on Outsourcing, which the industry is already familiar with. Similarly, MAS Notice FSMN21 Technology Risk Management (the "TRM Notice") focuses on identifying and maintaining high availability for "critical systems", the test for which is dependent on the consequences of the failure of that system. In both cases, heightened governance requirements are triggered by the materiality and impact of failure, rather than the breadth of usage or degree of operational integration. Also, the concept of "integration" may benefit from further clarity, as the current description in the Annex to the proposed AIRG may present practical challenges. It would also be difficult to measure the extent of integration by using the number of workflows in which the AI tool is used as an indicator because what constitutes a "workflow" is ambiguous and open to interpretation. Instead, precise, impact-based examples may be more useful in illustrating when FIs cross the threshold beyond which basic policies would not be sufficient. In addition, further clarity on the following would be appreciated: • How FIs should deal with third-party service providers using AI in their provision of services to FIs. Footnote 10 of proposed AIRG provides that third-party AI "includes all providers of thirdparty AI products and services, which can include third-party systems, models, as well as data used for AI. Existing third-party products and services where AI has been introduced would also be included". This would be resource-intensive for FIs to monitor from an operational perspective and to not only consistently identify AI across all relevant business and functional areas (potentially including a scenario where a service provider did not use AI when they were initially appointed but subsequently started to use AI in a manner that is not apparent to the FI), but also to establish and maintain an accurate and up-to-date inventory of all AI usage. • Whether service providers that do not necessarily provide an AI product and services but use AI in delivering those products and services are within the scope of third-party AI. Further, FIs may not be able to maintain oversight of third-party service providers’ use of AI to the degree that proposed AIRG require. For example, only a minority of FIs with greater bargaining power may be able to request contractual rights of disclosure/audit of the data used by such third-party service providers. This also means that FIs may not be able to adequately monitor existing third-party products and services for the introduction or removal of AI.

		Question 2: We understand that FIs would appreciate clarity on whether off-the-shelf tools with latent or autoenabled AI tools would be within the proposed scope – for example, word processing software potentially bundled with AI drafting tools but un-activated for the time being. The proposed AIRG appears to assume that whether a particular tool should be treated as AI is readily apparent and is a binary enquiry. We would appreciate further clarity on the following scenarios: (i) off-the-shelf software where AI functionality is optional and not in use by a party (e.g., where Microsoft Office is licensed, but Copilot is not enabled on an enterprise-wide basis); and (ii) off-the-shelf software where AI functionality is incidental to the performance of the primary function of the software (e.g., Microsoft Office is used primarily as a word processing programme and Copilot is embedded in the software to provide drafting assistance and other incidental functions; Slack is used primarily as a messaging platform but AI is used to improve search functionality) **Question 3:** Please refer to our observations in response to Question 1 above. **Question 4:** While the proposed expectation for FIs to establish a dedicated committee to oversee AI risk may be suitable for some FIs, we respectfully request more flexibility for FIs to build an oversight structure that is the most suitable for their particular circumstances, especially where FIs already have frameworks in place that can be adapted to address AI-specific risks. In addition, we also seek clarification on whether FIs may leverage group-wide resources when designating individuals to be responsible for overseeing AI risks **Question 5:** Please refer to our observations in response to Question 1 above. **Question 6:** Please refer to our observations iPlease refer to our observations in response to Questions 1 and 2 above. Our members have also highlighted that it is difficult for them to track the use of AI within an organisation to the degree contemplated by the proposed AIRG. Some members foresee practical issues with classifying whether certain AI use cases are "live" or decommissioned, especially where ad-hoc usage and "recommissioning" of an AI tool is standard practice. Further, even if employees are limited to approved versions of AI tools and such usage is accurately and timely monitored, FIs run the risk of failing to detect "shadow AI" usage where employees use opensource AI tools on their own initiative. We are also concerned with how detailed and prescriptive section 4 of the proposed AIRG is, and query whether how often or likely some of the expectations in section 4 would be applicable to most of our members. For example, paragraph 4.5 of the AIRG appears to

		address a specific scenario where a FI is commissioning a third party to develop an AI tool for the FI, which is a scenario that may not be applicable to all FIs, or in all circumstances. Even in a scenario where an FI is commissioning a third party to develop an AI tool for the FI, the FI is unlikely to be able to demand that the underlying data used by the third party is sufficiently representative or is of high quality, and may not even be able to assess the representativeness or quality of such underlying data if the developer does not disclose these to the FI.n response to Question 1 above. **Question 7:** Please refer to our observations in response to Question 1 above. **Question 8:** Our members have requested more clarity on why AI tools would require a wholly separate risk assessment and management process when it can simply be part of any existing business continuity or systemic risk assessment and management process. For example, an AI tool which comes bundled in a third-party tool should not require the FI to assess the AI separately regardless of what it can do or how it is used, given that the system as a whole is a single assessment. We also understand that FIs are concerned as to the practical feasibility of limiting AI usage as a possible measure of AI risk management. **Question 9:** Apart from the concerns raised in our responses to the questions above, we are also concerned that the extensive, granular requirements provided in the proposed AIRG may, if not carefully calibrated and sequenced, create practical frictions for further innovation, AI research and development, and responsible AI adoption, especially given the present trend towards deregulation in other key markets. For instance, the U.S. Commodity Futures Trading Commission is currently carrying out a pro-innovation framework while also reducing enforcement. Our members who have offices in multiple jurisdictions provided feedback that the expectations in the proposed AIRG are higher than the expectations in the jurisdictions that they operate in. Although proposed AIRG provide that FIs "may leverage AI risk management frameworks of their parent entities, as long as such frameworks meet the expectations set out in these Guidelines", our members highlighted that the relative demands of the proposed AIRG mean that any existing AI risk management policies that have been implemented on a group basis for compliance with existing requirements elsewhere may need to be revised so that they may comply with the higher standard provided in the proposed AIRG, and implemented in other jurisdictions as well. Alignment with global approaches is imperative – feedback from FIs with an international presence and operating across multiple locations, is that these measures in the Guidelines are relatively prescriptive, demanding and regulate ahead of other global financial markets and jurisdictions. We also observe that there is a substantial overlap between the proposed AIRG and the Technology Risk Management Guidelines. Given the overlap between the two sets of guidelines, including guidance specific to the management of AI risks as an addendum to

		the existing Technology Risk Management Guidelines may help to clarify incremental additional requirements and reduce the risk of unnecessary duplication. As mentioned in our response to Question 1 above, MAS's expectations in regards to managing AI risk can made be more consistent with existing guidelines relating to outsourcing and technology risk, leveraging on similar concepts already used in such existing guidelines. This would also afford FIs the flexibility to adapt existing frameworks to address AI risks, rather than having to create a new framework from scratch, which may prove to be duplicative when carried out in practice. **Question 10:** We respectfully request a longer transition time given the granular requirements provided in the proposed AIRG of 18 to 24 months. As mentioned in our response to question 7, FIs with offices in other jurisdictions may need to implement the higher AI risk management standards globally.
5	Andrew Koh M W	**Question 1:** Research in regulatory adaptation (Barth et al., 2006) supports proportionality as an effective mechanism to avoid undue compliance burdens while targeting systemic risks. I commend MAS for adopting this approach. To enhance clarity and operational utility, I propose the development of risk-tiered implementation pathways, which is informed by institutional size, AI adoption maturity, and customer impact thresholds, as seen in the EU's AI Act and the Bank of England's model risk frameworks. Proposed details as follows: Risk-Tiered Implementation Pathways for AI Governance ： 1. Core Principles · Proportionality: Requirements scale with institutional size and AI maturity, avoiding one-size-fits-all. · Legitimacy & Standards: Recognized global frameworks (EU AI Act, Bank of England's model risk). · Customer Impact Sensitivity: Higher thresholds of scrutiny where AI decisions materially affect consumer rights, financial stability, or systemic trust. 2. Tiering Dimensions Dimension Low-Risk Tier Moderate-Risk Tier High-Risk Tier Institutional Size Small fin-techs. Mid-sized banks, insurers Systemically important FIs AI Adoption Maturit Pilot-stage. Operational AI in core processes Enterprise-wide AI integration Customer Impact Minimal Moderate High 3. Operational Pathways · Low-Risk Entities

		o Simplified reporting templates o Light-touch supervisory reviews o Sandbox participation encouraged · Moderate-Risk Entities o Periodic model risk audits o Governance committees with AI oversight o Customer transparency disclosures · High-Risk Entities o Continuous monitoring and stress testing o Independent validation of AI models o Mandatory regulatory reporting and escalation protocols 4. Empirical Support · Studies on tiered regulatory adoption show reduced compliance ambiguity and smoother integration for non-bank financial institutions. · Evidence suggests graduated obligations foster innovation while maintaining systemic safeguards. · Tiering aligns with risk-based supervision models already proven in financial regulation. 5. Strategic Benefits · Clarity: Institutions know exactly what is expected at each stage. · Efficiency: Resources are allocated proportionally to risk. · Adoption: Non-bank entities face fewer barriers, encouraging responsible AI uptake. · Resilience: High-impact institutions subjected to rigorous oversight, protection of customers. **Question 2:** The proposed scope is appropriately broad and technology-agnostic, consistent with international AI governance trends (OECD, 2024). I recommend MAS to consider integrating an adaptive classification framework to distinguish between static automation and dynamic learning systems. Proposed Adaptive Classification Framework • Static automation: o Follows fixed rules, predictable outputs. o Risks: operational errors, rigidity, lack of adaptability. o Governance: compliance checks, deterministic audits. • Dynamic learning systems: o Continuously evolve with new data, sometimes in opaque ways. o Risks: bias drift, adversarial manipulation, explainability gaps. o Governance: ongoing monitoring, transparency standards, adaptive risk controls.

		Benefits for MAS • Regulatory clarity: MAS could set differentiated compliance requirements—lighter for static systems, stricter for dynamic ones. • Risk-based oversight: Supervisory intensity scales with system adaptivity and potential systemic impact. • Future-proofing: As AI adoption grows in finance, MAS avoids lumping all automation under outdated categories. Proposed Framework Elements • Classification tiers: o Tier 1: Static, rule-based automation. o Tier 2: Semi-adaptive systems (limited learning, bounded scope). o Tier 3: Fully dynamic learning systems (continuous adaptation, high uncertainty). • Assessment criteria: o Degree of autonomy. o Transparency/explainability. o Frequency of model updates. o Impact on financial stability and consumer protection. • Adaptive governance tools: o Continuous monitoring dashboards. o Mandatory model documentation and audit trails. o Stress-testing for dynamic systems. **Question 3:** Governance research (Cobbe et al., 2023) underscores the critical role of board-level AI literacy, not technical expertise, but functional understanding of AI's strategic and operational risks. I suggest incorporating explicit expectations for continuous AI governance training, potentially aligned with existing director certification programs. In addition, evidence from organizational studies supports the formalization of a Chief AI Risk Officer (CAIRO) role in larger Financial Institutions, with clear reporting lines to both senior management and the Board. **Question 4:** Literature on enterprise risk governance (Kaplan & Mikes, 2014) supports the establishment of dedicated cross-functional oversight bodies for emerging and complex risks. To assess organizational AI risk materiality, I suggest MAS could integrate a multi-factor scoring model drawing from dimensions such as: • Data criticality and volume • System autonomy levels • Regulatory and customer impact • Model interpretability and stability

		Dynamic reassessment intervals (e.g., semi-annual) should be mandated to reflect the evolving risk landscape. **Question 5:** IBM and governance researchers have highlighted the challenge of "shadow AI" in large organizations. I endorse the requirement for systematic AI identification, and suggests supplementing it with reference implementation frameworks, such as the NIST AI RMF (2023)'s "Map" function, to standardize identification protocols across FIs. Additionally, automated discovery tools using metadata and log analysis should be encouraged, as research indicates they improve detection of unsanctioned AI deployments. **Question 6:** An AI inventory is a foundational element of model governance (Bhatt et al., 2021). I recommend specifying core metadata fields aligned with international standards (e.g., ISO/IEC 42001) while allowing organizational customization. Periodic attestation mechanisms, modelled after SOX-like controls, could significantly strengthen accountability and ensure inventory accuracy over time. **Question 7:** The dimensions of Impact, Complexity, and Reliance align with widely cited risk taxonomies (Floridi et al., 2021). I propose augmenting these with a fourth dimension: Adaptability, which captures the degree to which an AI system can evolve post-deployment, a key risk factor for reinforcement learning and generative AI systems. Explainability should be treated as a critical sub-dimension of Complexity, given its direct link to regulatory compliance and operational trust. **Question 8:** The proposed controls are comprehensive and reflect a full lifecycle perspective, consistent with best practices in AI assurance (Raj et al., 2023). I recommend the inclusion of exemplar case studies, illustrating proportional control implementation, for instance, contrasting controls for a robo-advisor versus an internal document summarization tool. Further emphasis on continuous validation and real-time monitoring, particularly for generative AI systems, is supported by research on AI drift and performance decay (Liang et al., 2023). **Question 9:** Two areas merit further attention based on emerging research: 1. AI Supply Chain Governance: Extending third-party risk management to include open-source model provenance, training data lineage, and service-level AI integrity commitments. 2. AI Incident Disclosure Protocols: Structured reporting frameworks for AI failures, biases, or breaches, informed by cybersecurity incident disclosure regimes and aligned with international initiatives such as the G7 AI Principles.

		Question 10: Transition periods are widely used in regulatory rollouts to allow for organizational adaptation (Coglianese, 2022). A 12-month period is generally adequate but may be insufficient for smaller FIs with limited AI/risk management resources. I propose a staged compliance timeline, where: • Tier 1 Financial Institutions (high AI exposure) adhere to the full Guidelines within 12 months. • Tier 2/3 Financial Institutions are permitted an extended timeline (e.g., 18–24 months) for full lifecycle controls, provided they implement foundational policies and inventories within the first year.
6	Aon Singapore Pte Ltd	**Question 1:** No comments. **Question 2:** No comments. **Question 3:** Aon Singapore is part of the global Aon Group, which operates 500 offices in more than 120 countries. The Aon Group has adopted a globally consistent data and technology risk operating model, supported by our global AI Governance team. The team is trained and appropriately qualified to provide guidance and conduct risk assessments on AI-related activities and use cases. Therefore, it is not feasible to have a Singapore-only AI risk operating model nor have the Singapore Board and Senior Management directly perform some of the activities set out in Section 2: AI oversight. We propose that the (Singapore) Board and Senior Management of FIs that are branches or subsidiaries with parent entities in other jurisdictions be able to rely on their parent entities' global/ regional AI Governance team for Section 2 on AI oversight. For clarity, actions set out in Sections 2.3, 2.4, 2.5, 2.6 [excluding 2.6 (f)] can be completed on a global level, with regular assurance provided to Singapore Board and Senior Management to ensure effective oversight on AI risk management. Where gaps are identified between the FI's global AI governance standards and MAS' expectations, active internal discussions between the local Board and Senior Management or local personnel identified and the global AI Governance team will be taken to ensure alignment with MAS' standards, minimally for Singapore domiciled entities. Notwithstanding that the overall AI oversight is completed at the global level, the Singapore Board and Senior Management can be expected to be responsible for assessing personnel competence as outlined in Section 2.6 (f). **Question 4:**

		In line with our feedback regarding question 3, we suggest that FIs that are subsidiaries/ branches of global companies be exempted from this expectation to establish a dedicated local cross-functional committee to oversee AI risk if exposure is deemed material. Given that most of our AI technology is not solely for local use, and we have as global AI governance team which consists of representatives from different functions, we propose that FIs with a similar operating model can leverage on their global AI governance team as opposed to potentially duplicating such teams at the Singapore level. Nevertheless, at the local level, the representative identified in Section 2.6 (f) will be expected to review the AI risk exposure assessment. **Question 5:** No comments. **Question 6:** No comments. **Question 7:** May we seek clarification on whether Financial Institutions that already have AI governance frameworks and policies aligned with, or exceeding, the requirements set out in the EU AI Act (Regulation (EU) 2024/1689) may be exempted from duplicative compliance with the MAS guidelines? In particular, where an institution can demonstrate that its existing AI policies, controls, and documentation meet or surpass the safeguards envisaged under the MAS framework, would MAS consider: • Recognizing such EU AI Act–aligned arrangements as equivalent or substantially compliant, and • Allowing reliance on these existing policies in lieu of having to implement separate, MAS specific AI governance requirements? **Question 8:** We would like to highlight to MAS' attention of an element that may prove challenging to section 4.12 - the potential that third-party vendors engaged may not be forthcoming with information on the risks of their AI use case, system or model. Due to commercial reasons and the rapidly changing technology, vendors may choose to downplay limitations and bury risk disclosures when responding to such questions posed by FIs. It is also common for large global technology providers to not respond to RFPs or detailed risk-related disclosures. From a risk management perspective, it would be difficult for FIs to reach an objective conclusion when they are provided with selected information or minimal information. Nevertheless, we acknowledge that financial institutions should continue to be expected to request such information from developers, and to justify and document their model and vendor selection processes as far as reasonably possible, based on the information available to them at the time. **Question 9:**

		No comments. **Question 10:** No comments.
7	Asia Pacific Loan Market Association Limited	**Question 1:** 1. We welcome the MAS's robust and practical approach in applying the Guidelines proportionately across all FIs, regardless of size or business model. This ensures that smaller FIs are not unduly burdened while larger FIs remain accountable for more complex AI deployments. 2. We particularly support the flexibility for FIs operating in Singapore that are branches or subsidiaries of foreign parent entities to leverage their parent entities' AI risk management frameworks, provided that such frameworks meet the requirements set out in the Guidelines. This avoids duplication of efforts and promotes global consistency, while ensuring local compliance. 3. To further strengthen the proportionate application of the Guidelines, MAS may wish to expand the Annex with more examples of AI use cases, including those involving AI agents, given that some of the risks associated with AI agents are highlighted in paragraph 1.10 of the Guidelines. Such examples would help FIs align their proportionality analysis with the requirements of the Guidelines. **Question 2:** 1. We recognise that the proposed scope of AI use cases, systems and models set out in paragraph 1.2 of the Guidelines is intentionally broad given the rapidly evolving nature of AI applications. While this ensures comprehensive coverage, it does not distinguish between simple models (e.g. rule-based decision trees) and highly complex systems (e.g. deep learning models). 2. The MAS may consider providing additional guidance on tiering AI models by complexity and risk exposure, to further assist FIs in analysing the different types of AI applications that they use in their business and calibrating their AI risk management frameworks accordingly. 3. The examples of AI-related risks highlighted in paragraphs 1.8 to 1.10 are helpful. In addition, MAS may consider addressing risks arising from the automated and real-time nature of AI systems, which could exacerbate systemic failures. For example, digital withdrawals triggered by certain AI applications could accelerate a bank run much faster than traditional mechanisms. **Question 3:** We strongly support MAS' clear articulation of the Board and senior management's responsibilities in fostering appropriate risk culture and overseeing AI-related risks. This top-down accountability is crucial to embedding trust and resilience in AI adoption. **Question 4:** 1. We agree with MAS' proposal for FIs to establish a dedicated cross-functional committee where AI risk exposure is material. To ensure effectiveness, that committee should report

		directly to the Board, and not to a sub-committee of the Board, to avoid dilution of responsibility. 2. The MAS may consider highlighting that the cross-functional committee should be comprised of personnel with the relevant knowledge and experience (such as relating to AI, IT, data management or operational risk etc.) to ensure that technical and governance perspectives are integrated. **Question 5:** **Question 6:** **Question 7:** 1. We recognise that the AI risk materiality assessments may differ from FI to FI, and that the MAS has purposefully kept the Guidelines relatively practical and flexible to allow FIs certain degree of discretion to tailor their risk assessment criteria. 2. In addition to the proposed dimensions of impact, complexity and reliance, the MAS may consider including residual risk or uncontrollable risk (i.e. risk outside the FI's direct influence or risk that cannot be mitigated). 3. FIs may also conduct AI risk control (i.e. guardrails) assessments alongside materiality assessments. This dual approach would ensure not only that risks are analysed, but that mitigation strategies are also evaluated for effectiveness. **Question 8:** We endorse the MAS' proposal for FIs to implement robust standards and controls across the AI life cycle. To further strengthen this framework, FIs may also consider the assessment of the overall control risk, residual risk, and acceptance of any risk that cannot be mitigated. **Question 9:** We would recommend clarifying how these Guidelines interact with existing MAS guidelines on cybersecurity, data protection, operational resilience, outsourcing, and technology risk management etc. Cross-referencing relevant guidelines would help FIs adopt a holistic and comprehensive approach, as well as avoiding duplication. **Question 10:**
8	Asia Securities Industry & Financial Markets Association	**Question 1:** 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:** 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:** 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:** 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:** 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:** 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:** 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:** 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:** 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:** 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.
9	Asia Wealth Platform Pte. Ltd.	**Question 1:** With AI being an emerging technology that is rapidly maturing and providing significant benefits at scale, it is becoming increasingly difficult to track and control how AI is used across business use cases. Given this, leaving the decision to adopt an AI risk management framework to proportionality may not be ideal at this stage. The establishment of a formal AI Risk Management Framework should be mandated as a foundational requirement for all FIs. The proportionality should instead apply to how each FI chooses to operationalise and limit AI adoption, based on their defined risk appetite. **Question 2:** **Question 3:** **Question 4:** **Question 5:** **Question 6:** Inventorising the use of AI is only meaningful where it enables the monitoring and mitigation of risks that are not already addressed by existing controls. For example, while using tools such as Gemini for research technically constitutes AI usage, it does not necessarily introduce incremental risk when compared to traditional search engines **Question 7:**

		An additional risk component that should be considered relates to regulatory and legal risks arising from the implementation of AI. This includes potential legal liability stemming from AI-generated outputs, as well as risks of non-compliance with applicable laws, regulations. **Question 8:** **Question 9:** **Question 10:**
10	askOdin Pte Ltd	**Question 1:** We support the proportionate application based on materiality. However, we advocate that for AI models involved in capital allocation, credit decisions, and high-stakes due diligence, the standard for "Defensibility" must remain absolute, regardless of the FI's size. We propose that the Guidelines distinguish between "Probabilistic AI" (Generative content creation) and "Deterministic AI" (Logic verification). FIs using Generative AI for investment judgment should be required to implement a "Verification Layer" (compile-time logic checks) to ensure that narrative claims extracted from unstructured data are mathematically consistent before human review. Proportionate application should not excuse the lack of deterministic guardrails in financial decision-making. **Question 2:** The scope should explicitly include "AI for Unstructured Data Analysis in Due Diligence and Investment Research." Currently, many FIs use LLMs to summarize pitch decks, earnings calls, and legal documents. This creates a hidden "Hallucination Risk" where financial facts are distorted by the model's probabilistic nature. The Guidelines should categorize any AI used to interpret financial narratives or extract structured data from unstructured sources as a "Material Risk" system. Such systems must require forensic provenance tracking (the ability to link a specific output back to the specific source text character offset). **Question 3:** The Board should not just oversee "Policy," but "Architecture." We recommend that Senior Management be required to demonstrate that their AI systems possess "Traceability"—the ability to audit why a model reached a conclusion. If an AI recommends an investment or flags a risk, the Board must be able to see the specific data lineage (Source Document → Extraction → Logic Check → Decision). "Black Box" recommendations, where the reasoning path cannot be audited step-by-step, should be deemed unacceptable for fiduciary duties. **Question 4:** **Question 5:**

		Question 6: **Question 7:** In addition to Impact, Complexity, and Reliance, MAS should include "Semantic Brittleness" as a risk dimension for Generative AI. Semantic Brittleness refers to the risk of AI models misinterpreting vague qualitative statements (e.g., "significant growth") as precise quantitative facts. Risk assessments should measure a model's ability to distinguish between Fact (auditable data) and Narrative (founder claims). Systems that cannot mathematically verify the logic of a narrative pose a solvency risk. **Question 8:** **Question 9:** Topic: The Systemic Risk of Agentic AI and the Need for Deterministic Verification As the financial sector moves toward "Agentic AI" (systems that not only analyze but act), reliance on purely probabilistic models (LLMs) for capital allocation tasks creates a new class of systemic risk. We strongly advocate that the Guidelines encourage the adoption of "Hybrid Architectures." For high-stakes financial use cases (Credit, Investment, Solvency), natural language should be converted into an Intermediate Representation (IR)—executable, logic-validated code—before any decision is rendered. This "Financial Compilation" approach ensures that the reasoning process is mathematically consistent (a compile-time check) before it is acted upon. This prevents hallucinations at the structural level and provides the granular audit trail necessary for effective regulatory supervision. **Question 10:**
11	AXA XL which includes "Catlin Singapore Pte Ltd", "XL Insurance Company SE (Singapore) Branch" and "XL Re Europe SE (Singapore) Branch"	**Question 1:** We consider that the proportionate approach as set out in set out in paragraph 1.5 of the proposed Guidelines is a fair approach supporting effective and tailored risk management. This approach is consistent with how large FIs structure AI risk management frameworks. For example, AXA XL entities are subject to an enterprise-wide AI Policy, which establishes a risk based and proportionate AI Risk Management framework across the AI lifecycle to identify, assess and manage AI related risks. **Question 2:** We note that the definition of AI in the Guidelines is similar to that set out in the EU AI Act. We consider that the proposed scope of AI for the purposes of the Guidelines is reasonable and is aligned with AXA XL Policy. **Question 3:**

		We consider that the proposed responsibilities of the Board and senior management in overseeing AI risk management are reasonable and generally aligned with AXA XL AI Governance Framework and Policy. **Question 4:** We have no specific comments on this proposal but will note that AXA XL's existing AI Governance process would fulfill these expectations by integrating all relevant business functions and enabling informed decision making at key lifecycle changes. **Question 5:** AXA XL supports MAS' proposal to implement clear definitions, identification criteria and processes for consistent identification of AI use across an FI. AXA XL AI definitions of AI, AI Systems and Models is aligned with proposal. **Question 6:** AXA XL supports the proposal to establish and maintain an accurate and up-to date inventory of all AI usage. This requirement is captured within the AXA XL AI Policy. **Question 7:** AXA XL is generally comfortable with this proposal. We note that whilst the terminology used is different to the equivalent risk assessment in EU legislation, the underlying principles appear to be the same. The EU AI Act is prescriptive legislation and sets out what are "prohibited AI practices" and "high risk AI systems". Although we appreciate that Guidelines are more flexible for FIs, as the Guidelines leave the risk assessment to the FIs and only provide various risk dimensions (impact, complexity, reliance) to be considered, different FIs could assess the same/similar AI system differently. As the Guidelines specify that "Such assessments are critical for calibrating the AI risk management approach, ensuring that controls are applied proportionate to the risks posed by AI. This ensures that the AI use case, system or model assessed as having higher risk materiality receives more stringent scrutiny and robust controls, supporting effective oversight and confirming that AI usage remains within approved boundaries". This could result in different controls being applied to the same AI system by different FIs. For certainty AXA XL request that the Guidance provides examples of what would constitute material/high-risk AI systems. **Question 8:** AXA XL supports the proposed AI life cycle controls. AXA XL's AI Governance Framework already embeds equivalent controls across the AI lifecycle. **Question 9:** None. **Question 10:** AXA XL supports the proposed 12 month transition period.
12	AXS Pte Ltd	**Question 1:**

		We recommend that the Guidelines include clearer, illustrative examples to differentiate between "basic" and "integrated" applications of AI. Such clarification would help prevent unnecessary "compliance creep," where smaller payment institutions may feel obliged to adopt enterprise level AI governance frameworks even when deploying low risk, supportive AI tools **Question 2:** With respect to paragraph (b), which defines a "System" as comprising models and "other machine-based components," we request clarification on the intended scope of these components. Specifically, does this encompass hardware infrastructure, data pipelines (ETL processes), or user interfaces? For compliance purposes, it is critical to delineate the "boundary" of an AI system to ensure firms can maintain an accurate AI Inventory as required under Section 2. We recommend that MAS limit the definition to components that directly influence the model's output or its performance monitoring, thereby avoiding unnecessary inclusion of peripheral systems that do not materially affect AI risk. Regarding paragraph (d), which notes that AI may "vary in their levels of autonomy and adaptiveness after deployment," we seek clarification on whether static machine learning models, those trained once and not designed to update or adapt parameters in the production environment, are considered within scope. Clear guidance on this point would help firms distinguish between adaptive AI systems and static models, ensuring proportional governance measures are applied. **Question 3:** **Question 4:** We recommend that MAS amend the guidance to clarify that financial institutions may leverage existing governance structures to satisfy the requirement for a cross functional committee, provided they can demonstrate that AI specific risks are being addressed with appropriate expertise and regularity. Establishing a standalone "AI Committee" risks fragmenting decision making and fostering a siloed view of risk, where AI is treated as separate from the broader technological ecosystem. For small and medium sized institutions, the administrative burden of maintaining an additional committee could contribute to "committee fatigue" and dilute focus on core risk issues. We therefore suggest that MAS permit the "dedicated" requirement to be met through a formalized sub mandate or a standing agenda item within an existing high level risk committee. This approach would ensure proportionality, integration, and effective oversight without imposing unnecessary governance overhead **Question 5:** **Question 6:** **Question 7:** **Question 8:**

		Question 9: **Question 10:** We recommend extending the proposed transition period from 12 months to 18 months, in light of the following challenges: * Many existing models may not currently meet the proposed standards for explainability or transparency. Financial institutions will require sufficient time to either uplift these models or safely decommission and replace them, without risking disruption to critical payment flows. * The current market has a limited pool of independent AI auditors and validation experts. Compressing the transition into 12 months could create a bottleneck, resulting in lower quality validations or significantly inflated compliance costs, particularly for smaller institutions. An 18 month transition period would provide firms with the necessary runway to implement changes responsibly, while ensuring that compliance efforts remain both effective and proportionate.
13	Behavox Ltd	**Question 1:** Behavox strongly supports MAS's proportionate application framework, including the distinction between AI used as an integrated part of business processes and lower-risk, assistive AI. The tiered approach is practical and appropriately focuses resources and controls where risk is greatest, while maintaining baseline governance across all FIs. To support consistent application across the industry, additional clarification would be helpful in three areas: Borderline "integrated use" scenarios: Guidance or examples addressing cases where AI is increasingly relied upon across teams or embedded in workflows but retains meaningful human review would assist consistent tier classification. Minimum baseline documentation for low-risk tools: While lighter oversight is appropriate, minimal expectations that preserve an audit trail (e.g., approved tool register, user guidance, and periodic attestation to baseline policies) would help FIs demonstrate control without creating undue burden. Transition triggers and reassessment: As AI adoption grows incrementally, practical guidance on when FIs should reassess tier classification (e.g., periodic reassessment, changes in reliance, or workflow dependency) would reduce ambiguity and promote timely uplift of controls. **Question 2:** Behavox supports MAS's comprehensive scope, including explicit coverage of machine learning, Generative AI, and AI agents. This inclusiveness helps prevent regulatory gaps as technology evolves. The exclusion of purely deterministic calculators and tools is also appropriate.

		In practice, financial institutions deploy a range of materially different AI approaches, which present distinct risk characteristics and supervisory considerations. For example, Generative AI chat interfaces, classification and scoring models used for detection or prioritisation, and reinforcement or agentic systems that adapt behaviour over time differ significantly in determinism, explainability, testing methods, and change dynamics. Treating these approaches uniformly may lead to inconsistent outcomes in validation, monitoring, and change management. Behavox therefore encourages guidance that recognises these distinctions, enabling proportional and risk-appropriate application of controls across different AI classes while preserving the core principles of AIRG. To improve consistency in application, Behavox suggests further clarification on: Hybrid systems: Many real-world implementations combine rule-based logic with machine learning (e.g., rules-based filtering followed by ML risk scoring; rule-based chatbot flows with LLM escalation). Additional guidance on when the AI component becomes sufficiently material to warrant governance of the overall system, versus governance of discrete AI components, would be valuable. Documentation granularity (model vs system vs use case): AIRG frequently references "AI model, system or use case." Guidance on when system-level documentation is sufficient, versus when model-level documentation is required (particularly where constituent models have different risk profiles or can be used independently), would improve scalability. Development, pilot, and production status: Clarifying expectations for AI under development or in pilot, including how such systems should be identified, tracked, and governed prior to full production deployment, would support disciplined lifecycle control. **Question 3:** Behavox supports MAS's emphasis on Board and senior management accountability for AI risk management, including alignment with enterprise risk management and clear escalation and reporting. This is consistent with international good practice. As AI adoption matures across regulated institutions, effective Board and senior management oversight of AI risk depends on clear reporting that distinguishes between materially different AI deployments, as oversight expectations and escalation thresholds vary by AI type and use case. Given the prevalence of third-party AI in the FI ecosystem, additional guidance would help clarify expectations around: Provider-facing governance artefacts supporting FI due diligence: FIs frequently require documentation and evidence from providers to meet their own governance obligations. MAS could consider clarifying the minimum categories of provider information that would typically support FI oversight (e.g., governance framework summary, validation approach, key limitations, monitoring approach, and change notification practices), while recognising confidentiality and intellectual property constraints.

		Harmonisation across global frameworks: Many FIs and major providers operate globally. Guidance on how AIRG expectations relate to internationally recognised standards (e.g., ISO/IEC 42001 and established model risk management principles) would improve implementation efficiency and reduce unnecessary duplication. Behavox notes that FI accountability for safe and responsible use remains paramount; clearer expectations for provider inputs can help FIs discharge accountability more effectively. **Question 4:** Behavox supports MAS's proposal. Enterprise-level coordination is essential where AI use is widespread across functions, where reliance is high, or where AI supports critical regulated activities. In Behavox's experience building and operating AI-native controls platforms, effective enterprise-level oversight does not emerge organically from organisational structure alone, it must be deliberately supported by how AI systems are designed and deployed. Behavox has therefore built its platform to operate across compliance, risk, legal, technology, and audit functions, with shared data and consistent controls. This cross-functional design enables institutions to view AI risk holistically, rather than through fragmented functional lenses, and supports clearer ownership, escalation, and challenge as reliance on AI increases across the organisation. To support consistent assessment of overall AI risk exposure, MAS could consider providing examples of practical aggregation approaches, such as: How to consider concentration of high-materiality use cases across business lines How to account for a large number of medium-risk AI use cases that collectively increase exposure Escalation triggers tied to changing portfolio composition (e.g., increasing dependence on AI in critical workflows) These examples could be included in supplementary materials or the Project MindForge handbook. **Question 5:** Behavox agrees that AI identification is foundational. Clear definitions, designated ownership, and robust supporting processes are necessary to ensure visibility and control, including detection of unapproved or "shadow" AI use. For example, generative AI tools used to support investigation or analyst review are often deployed as user-facing capabilities, while classification and scoring models embedded within surveillance or risk workflows operate continuously and at scale in the background. More adaptive or agentic components may be orchestrated across multiple systems and functions. Without identification criteria that recognise these differences, firms risk AI inventories that are either incomplete or overly broad, making them harder to govern, validate, and oversee as the number and reliance on AI deployments grows.

		Two areas where practical guidance may further support implementation: Hybrid and embedded AI identification: Where AI components are embedded within broader systems (including vendor platforms), guidance on how FIs should identify AI at appropriate granularity would reduce inconsistency. Time-limited and experimental use: Clarifying how pilots, proofs-of-concept, or time-limited deployments should be captured and governed, including expectations for periodic attestation and review, would help FIs manage risk without stifling controlled experimentation. **Question 6:** Behavox supports MAS's inventory requirement and recognises its importance for accountability, oversight, and portfolio-level risk management. Our experience building and operating AI-led controls across multiple products shows that maintaining an accurate AI inventory becomes increasingly challenging as AI use scales across systems and functions. Inventories are most effective when they are closely tied to how AI models are configured, deployed, monitored, and changed in practice, rather than treated as a static register maintained in isolation. This is particularly important where AI is embedded within broader platforms or workflows, and where responsibilities for development, deployment, and oversight are shared between providers and financial institutions. Behavox suggests MAS consider further clarification on: Minimum mandatory vs recommended attributes: Differentiating baseline inventory fields from additional attributes expected for higher-materiality AI would support proportionality and audit consistency. Provider-to-FI inventory information sharing: Where AI is sourced externally, a practical model is a two-layer inventory: Provider-maintained inventory elements relating to model/system development, intended use, known limitations, and change history. FI-maintained deployment inventory elements capturing local configuration, use-case context, integration points, materiality rating, validation outcomes, and ongoing monitoring in the FI environment. Clarifying expected provider inputs (without shifting accountability away from FIs) would support efficient due diligence and reduce duplication. Triggers for inventory updates: Clear examples of "material change" events (e.g., retraining, intended-use changes, major version updates, or changes to reliance/human oversight arrangements) would help FIs maintain accurate inventories without creating excessive administrative burden.

		Question 7: Behavox supports MAS's three-dimensional approach. It provides a practical basis to calibrate controls and investment relative to risk and aligns with global approaches. In building and operating AI-led control systems, Behavox applies similar risk dimensions through alignment with international frameworks such as ISO/IEC 42001 and emerging EU AI Act–style risk categorisations. These dimensions are treated as dynamic rather than static, with continuous monitoring used to reassess materiality as models, use cases, and reliance evolve over time. This approach reinforces the importance of reassessing residual risk after controls, rather than relying solely on point-in-time classification. Behavox suggests additional operational clarity in three areas: Granularity of scoring: Whether MAS expects qualitative ratings (e.g., low/medium/high) or permits quantitative scoring methods, and what "good practice" looks like for consistent application. Generative AI hallucination risk: Hallucination and non-deterministic failure modes can amplify downstream impact even where reliance is limited. Guidance on how to reflect hallucination propensity and mitigation strategies within impact/complexity/reliance assessments would support consistency. Human oversight and reliance calibration: Reliance should reflect not only whether a human is "in the loop," but whether oversight is meaningful and auditable, including clear authority to override, training to counter automation bias, and evidence that overrides occur where warranted. A simple descriptive taxonomy (AI assists → AI recommends → AI decides with audit → autonomous) can help FIs calibrate reliance consistently. Behavox also supports MAS's emphasis on assessing inherent risk prior to controls and residual risk after controls, and notes that guidance on how best to document this distinction would help drive consistent risk acceptance decisions. **Question 8:** Behavox supports MAS's lifecycle-centric framework and agrees that these controls are essential for safe, fair, secure, and reliable AI. In practice, effective lifecycle governance depends not only on the existence of controls, but on the availability of clear, consistent documentation that allows those controls to be understood, tested, challenged, and maintained over time by financial institutions, providers, and supervisors. Documentation is the means by which assumptions, limitations, validation outcomes, and change decisions are made transparent across the AI lifecycle, including where responsibilities are shared between AI providers and deploying FIs. Behavox strongly recommends that MAS explicitly require comprehensive documentation for all AI deployments to support model risk management, internal audit, and regulatory oversight. As an AI provider supporting a broad range of regulated financial institutions, Behavox designs its AI systems and governance processes to meet the needs of model risk management, internal audit, and supervisory review. Behavox's model risk documentation has been shaped through direct feedback from model risk teams and regulators in the US, the UK, and other jurisdictions, and is continuously refined to remain current and as

		governance expectations evolve. Behavox would welcome the opportunity to share these lessons learned with MAS as AIRG implementation guidance continues to evolve. To support consistent and proportionate implementation across different AI types and risk tiers, Behavox highlights the following areas where additional operational guidance would be beneficial: Third-party and black-box AI compensatory testing: Where foundational or third-party models cannot be fully inspected, MAS could clarify what constitutes sufficient compensatory testing and documentation, and how expectations should be calibrated by risk materiality. Clear framing of shared responsibilities (provider generic testing and transparency; FI deployment-specific testing and oversight) would reduce gaps and inefficiency. Explainability proportionality by AI type: Practical examples or illustrative standards across AI classes (e.g., deep learning classifiers, behavioural surveillance models, Generative AI tools, assistive copilots) would help FIs and providers implement explainability that is proportionate, useful to investigators and decision-makers, and defensible for audit and regulatory scrutiny. Training data governance, IP rights, and data preparation: Additional guidance would be helpful on expectations for the use of historical real-world data in AI training, including data cleansing, anonymisation, aggregation, and intellectual property considerations. In practice, AI models are often developed using a combination of client-derived data (subject to contractual and legal constraints), cross-client aggregated insights, and synthetic data. Clear supervisory expectations on acceptable preparation standards and how firms should evidence lawful, ethical, and IP-compliant data use would support consistent lifecycle governance and auditability. Validation independence: Behavox strongly recommends that all AI deployments undergo robust independent testing to verify that actual performance aligns with the performance characteristics declared by AI developers and providers. This is critical to prevent over-reliance on vendor claims and to ensure FIs have objective evidence that AI systems operate as intended in their specific deployment context. Independent validators should possess the necessary separation, expertise, and authority to challenge development teams, particularly for high-materiality AI where outcomes may influence regulated decisions or controls. Validators must have the ability to test AI performance objectively under realistic conditions and to prevent deployment where risks are not adequately addressed or where declared performance cannot be substantiated through independent testing. From Behavox's experience deploying AI in regulated production environments, the most effective validation processes involve independent teams that: Test AI using realistic data representative of deployment conditions Verify performance claims through their own testing methodology Have authority to challenge or reject AI that fails to meet standards Document assumptions, limitations, and validation outcomes comprehensively

		MAS guidance on minimum independence standards and validation requirements would support consistent practice across the industry. Monitoring, drift, and change management: Guidance on calibrating monitoring intensity, drift indicators, and change triggers based on materiality would be particularly helpful for dynamic AI systems and frequently updated models. For mission-critical applications such as risk management, compliance surveillance, and customer advisory, Behavox strongly recommends that all AI deployments be locked in production, with any changes requiring formal change management, independent re-validation, and approval. This locked-model approach provides firms with greater control, prevents unintended model drift, and ensures that production performance remains aligned with approved testing outcomes. **Question 9:** Behavox commends the breadth of AIRG and highlights three areas that may warrant supplementary guidance as practices mature: Provider–deployer responsibility allocation: AIRG correctly emphasises FI accountability for AI risks, including third-party AI. However, clearer guidance on how responsibilities are commonly allocated across the AI supply chain (e.g., provider development controls and generic validation; FI deployment validation, use-case appropriateness, and oversight) would reduce the risk of both gaps and duplicative effort. A responsibility-allocation framework could support consistent, effective outcomes without diluting FI accountability. Synthetic data use: Synthetic data is increasingly used for testing, evaluation, fairness assessment, and scenario coverage. Supplementary guidance on assessing representativeness, privacy risks, and fitness-for-purpose, and how synthetic data practices should align with broader data governance expectations, would support safe adoption. Cross-border AI operations: Many FIs rely on global infrastructure and providers. Additional guidance on cross-border inference, logging, monitoring, auditability, and handling potential conflicts between jurisdictional requirements would support consistent compliance and robust risk management. Behavox notes these topics may be suitable for inclusion in supplementary materials (e.g., the Project MindForge handbook) rather than the initial Guidelines. **Question 10:** Behavox supports MAS's proposed 12-month transition period and agrees that a clear endpoint supports timely uplift of governance practices. Given practical constraints, Behavox recommends MAS consider a risk-based phasing approach, prioritising high-materiality AI systems for earlier compliance, consistent with the proportionality philosophy of AIRG. In addition, many FIs rely on third-party providers for key artefacts and operational inputs (e.g., documentation, validation summaries, monitoring and change notifications). Recognising this dependency in transition

		expectations would support achievable implementation while maintaining safety and accountability, including through compensatory controls where needed. Conclusion Behavox commends MAS for developing comprehensive and forward-looking Guidelines that recognise both the opportunities and the real operational risks associated with AI in financial services. AIRG represents a meaningful step toward consistent, lifecycle-based AI governance in financial services. From extensive experience deploying AI in regulated production environments, Behavox supports the overall direction of the framework and has highlighted a small number of areas where additional operational clarity would help firms and providers implement AIRG in a way that is both effective and proportionate. These include the governance of hybrid and foundation model architectures, calibration of explainability and oversight across different AI types and risk tiers, practical tiering as reliance on AI increases, and clearer treatment of shared responsibilities across provider–deployer ecosystems. Behavox stands ready to support MAS-regulated institutions in meeting AIRG expectations and to contribute constructively to industry efforts, including the Project MindForge AI Risk Management Handbook and related working groups. Drawing on hands-on experience with model governance, validation, and change management, Behavox would welcome continued engagement as AIRG is translated from principles into day-to-day supervisory and operational practice.
14	Blitz Consulting & Coaching	**Question 1:** The proportionality approach is appropriate and necessary. FIs in Singapore vary widely in product complexity and internal AI capability. Hence, applying the guidelines uniformly without proportionality would create uneven compliance burdens on the FIs. A proportionate framework would enable : * High-reliance AI systems receive deeper oversight * Smaller institutions adopting low-complexity AI can comply with fewer operational costs To help FIs gauge the proportionality, we suggest that MAS could consider adding short illustrative examples of the use-cases, e.g., a credit scoring model vs a conversational chatbot **Question 2:** The proposed scope is broad enough. One key suggestion is to clarify that "AI systems" should include hybrid pipelines where rule-based, ML (basic machine learning), DL(deep learning), and LLM (large language models) components may co-exist. Many real-world deployments have such pipelines, where the risk may arise from the interaction of these components instead of a single AI model. **Question 3:** MAS's emphasis on oversight rather than technical intervention is correctly aligned. Some additional clarifications, as follows, could strengthen implementation consistency: * Boards must ensure a named owner for each material AI use case * Execution must be delegated to management, with the boards focusing on Policy and Oversight.

		Question 4: AI risk cannot be adequately governed in silo's with only IT or Audit, or Product handling it. In terms of interpreting the extent of the exposure, some parameters along the line of "proposed risk dimensions" could be : *Impact: Potential customer or end-user harm; financial exposure; or operational downtime/disruption *Reliance: the degree to which decisions depend on AI output or autonomous behaviour *Complexity: black-box effect or lack of explainability in workflows; difficulty in tracing or validating, or explaining model behaviour in understandable language Institutions should be allowed to use their internal benchmarks as long as they can logically explain the reasoning behind those. **Question 5:** It will be helpful if there were a clear definition of the "AI lifecycle"; this will support a common vocabulary across functions such as Compliance, Risk, IT, and Model validation teams. **Question 6:** Some practical aspects to be considered: * FIs need flexibility to maintain a single combined inventory for AI + traditional models * Real-time updates may not be feasible for many FIs. Instead, MAS could consider Periodic updates with suitable disincentives for non-compliance. **Question 7:** These are pragmatic and essential. They form a useful mechanism for the determination of the extent of oversight needed, along with documentation levels. **Question 8:** The lifecycle framework is very robust; as such, it does not need additions. A pertinent aspect is differentiating high-risk AI systems from low-risk AI systems to support the adoption of simpler systems and reinforcing proportionality. **Question 9:** Some optional guidelines which MAS can consider: * Documentation expected for different risk levels * Handling of hybrid workflows (AI combined with decision rules or HITL/ human overriding AI) **Question 10:** A 12-month transition period is reasonable given the pace at which AI is evolving. Larger organisations with validation processes in place will be able to adapt quickly. Smaller organisations may require additional time for capability building, defining compliant processes, setting up an AI inventory, and developing templates.
15	CFA Society Singapore	**Question 1:** We support proportionate application to all FIs, consistent with MAS' risk-based regulatory approach.

		• Proportionality should be driven primarily by use-case materiality: (i) impact on customer outcomes, fairness and trust; (ii) complexity; and (iii) reliance/autonomy in decision-making. • The Guidelines should clearly differentiate AI that makes or materially shapes decisions (e.g., recommendation ranking, underwriting decisions, surveillance triggers) versus AI that assists humans (e.g., summarisation/drafting) – with stricter controls for the former. • For AI used in regulated activities (financial advice, suitability, portfolio construction, discretionary decision support), MAS could specify a minimum control floor regardless of firm size, to protect clients and avoid uneven standards. • We recommend MAS include illustrative, tiered examples in the Annex (e.g., low-impact internal productivity tools; medium-impact customer communications; high-impact advisory/portfolio and underwriting use cases) to reduce interpretation variance across firms. **Question 2:** We broadly agree with the proposed scope covering AI use cases, systems and models, including ML/DL, Generative AI and AI agents. • We recommend MAS explicitly address AI embedded within third-party tools and enterprise software (e.g., CRM, productivity suites, vendor advisory tooling) – including how Singapore entities should manage risk where models are developed/hosted offshore but used locally. • Composite advisory architectures are common: a deterministic suitability/product-governance layer combined with probabilistic/GenAI components for explanation or drafting. It would be helpful for MAS to recognise and encourage such constrained AI patterns where appropriate. • MAS should clarify expectations for distinguishing firm-approved AI use versus employees' personal AI use at work, and reinforce that employment/IT policies should prohibit or tightly restrict unauthorised personal/public AI tools on corporate systems. **Question 3:** We support MAS' emphasis on Board and senior management oversight, aligned with existing technology risk governance expectations. • We noted a potential gap in technical understanding at senior levels. Meaningful oversight requires baseline AI competency (training and/or access to expert support) for Board and senior leaders. • Where AI materially affects customer outcomes (e.g., advisory, suitability, underwriting), senior management should be accountable for calibration, monitoring and escalation processes, including clear decision rights and kill switch authority for severe incidents.

		• We recommend at least annual reviews of AI policies, inventories, materiality assessments and controls, with reporting to the Board on high-risk use cases and key risk indicators (KRIs). • Board reporting for advisory-related AI could include outcome indicators (e.g., suitability breach rates, override/challenge rates, complaint themes, drift/instability metrics, and incident trends). **Question 4:** We support establishing a dedicated cross-functional committee where overall AI risk exposure is material. • For high-impact use cases, membership should include technology and data experts, model risk management, compliance/legal, business owners and frontline users (e.g., advisory/wealth, credit) to ensure controls are practical and effective. • The committee should oversee shadow AI detection and remediation, and ensure that procurement, pre-production approvals and change management gates are consistently applied across business units. • Overall AI risk exposure should consider both the number/complexity of models and the scale of customer/business exposure (e.g., assets, customer segments, transaction volumes) affected by AI-enabled processes. **Question 5:** We support the proposal for clear definitions, criteria and processes to identify AI usage across the FI. • Identification processes should explicitly cover unauthorised AI tools and locally built shadow models (e.g., spreadsheets, scripts, prompt workflows) that can materially affect outcomes without formal governance. • Firms should tag AI by decision role (decisioning vs assistive), level of autonomy/reliance, and whether outputs are deterministic vs probabilistic/generative. • Controls should also address employee use of personal/public AI tools on corporate devices and networks through IT restrictions and HR policies, to reduce data leakage and compliance blind spots. **Question 6:** We support maintaining an accurate and up-to-date AI inventory, and recommend it be impact-focused rather than a mere list of tools. • Inventories should capture use-case purpose, business owners, affected customer segments/processes, and how AI outputs influence decisions and communications.

		• Include materiality ratings (impact, complexity, reliance/autonomy) with rationale; key dependencies (data sources, systems, vendors); and the update mechanism (model/version changes, prompt/template changes, auto-updates). • For client-facing or outcome-shaping use cases, inventories should link to validation evidence, monitoring dashboards, incident logs and audit trails to support supervisory review. • Where AI is embedded in third-party services, inventories should capture vendor accountability arrangements, due diligence status, and contingency/exit plans. **Question 7:** We agree that impact, complexity and reliance are strong core dimensions for AI materiality assessments. • For advisory/suitability contexts, we recommend MAS clarify expectations on reproducibility and auditability: reproducibility at the category/constraint level (e.g., eligible universe, suitability constraints, rationale factors) is a practical standard, even if narrative wording varies. • We recommend explicit consideration of (i) fairness/non-disadvantage risk; (ii) determinism/stability over time for identical inputs; (iii) client vulnerability and harm severity; (iv) effective reliance (actual override/challenge behaviour); and (v) and third-party opacity (prompt injection, data poisoning, leakage). • For GenAI in particular, consider truthfulness and misleading-risk (e.g., hallucinations or overconfident framing) as a driver of harm and trust erosion, especially where customers or frontline staff may over-rely on outputs. • Materiality assessments should also consider whether AI use may affect market integrity and public trust (e.g., misleading content, deepfake or manipulation risks). **Question 8:** We strongly support the proposed lifecycle controls across inception, development, validation, deployment, monitoring, change management and decommissioning. For high-impact use cases, controls should be calibrated to ensure suitability, auditability and security are not compromised by probabilistic behaviour. 8.1 Suitability and advisory controls: deterministic constraints with constrained GenAI • For advisory/portfolio workflows, firms should calibrate AI to be more deterministic where necessary to meet suitability obligations (e.g., hard stops on risk mismatch, concentration, product eligibility). • GenAI components should be constrained to pre-approved eligible universes and grounded on controlled data sources; narrative variability may be acceptable if the underlying constraints and reasoning factors are reproducible and logged. • Controls should explicitly address GenAI failure modes such as hallucinations and overconfident framing (e.g., grounding requirements, refusal behaviours, human review for client-facing text, and clear confidence/limitations cues).

		8.2 Reproducibility and auditability: decision lineage • Auditability should include full decision lineage: client inputs used, deterministic rule outcomes, model/version identifiers, prompts/templates, retrieved sources, and timestamps. • Firms should define practical standards for category-level reproducibility (constraints and rationale factors) and maintain logs sufficient to reconstruct why a recommendation was made and what controls applied. 8.3 Shadow AI, employee controls and data leakage prevention • Operational controls should proactively detect unauthorised AI usage (shadow AI) and enforce approved tooling via procurement and IT controls. • Employee use of personal/public AI tools on corporate networks should be restricted through policy, endpoint controls and monitoring, given data confidentiality and leakage risks. 8.4 Third-party AI vendor management and accountability • Where firms rely on third-party AI, due diligence should assess accuracy, robustness, security, and transparency; contractual clauses should address accountability, incident notification, audit rights and change management. • MAS could clarify how AI vendor management should integrate with existing TRM and outsourcing guidelines, including expectations where liability transfer may be limited in practice. 8.5 Human oversight and over-reliance risks • Advisors remain ultimately responsible for recommendations except in the cases of AI based robo-advisories; governance should include override/challenge expectations and monitoring for automation bias (e.g., unusually low challenge rates). • Firms should design UI and workflows to encourage critical review, and establish escalation routes where AI outputs conflict with policy or client interests. 8.6 Transparency and client disclosure • Where AI materially influences advice or customer outcomes, disclosure to clients regarding AI's role and limitations can improve transparency and market-driven accountability, while preserving the FI's responsibility for outcomes. **Question 9:** Overall, the Guidelines are comprehensive. We recommend additional clarifications in three areas: • Advisory and suitability: include a short consolidated sub-section (or Annex examples) for AI in financial advisory/portfolio management, covering suitability constraints, category-level reproducibility expectations, and audit lineage. • Third-party embedded AI and role clarity: clarify expectations for AI embedded in vendor tools and the respective responsibilities of manufacturers versus distributors of AI-enabled services in financial contexts.

		• Vendor accountability and outsourcing alignment: clarify how accountability, due diligence and ongoing monitoring expectations interact with outsourcing requirements and cross-border group governance. • IP and innovation: maintain flexibility and avoid approaches that force disclosure of proprietary methods; focus on governance, monitoring and outcomes rather than requiring detailed model disclosure. • Regulatory evolution: consider a periodic review mechanism (e.g., half-yearly or annual) to update expectations as AI transparency and stability improve, potentially shifting from prescriptive controls to more principle-based guidance over time. • Cross-jurisdictional alignment: benchmark with international regulators where appropriate to reduce fragmentation and implementation burden for globally operating firms. **Question 10:** We agree that a 12-month transition period is reasonable, provided implementation is prioritised by materiality and client impact. • We suggest MAS encourage interim milestones for high-impact use cases (e.g., within the first 3–6 months: complete AI identification and impact-focused inventory; classify use cases by materiality; implement minimum safeguards for advisory/suitability and other outcome-shaping uses). • MAS could also signal an intention to re-evaluate the Guidelines on a scheduled basis (e.g., within 6–12 months after issuance) to incorporate lessons learned and evolving risks, particularly for GenAI and agents.
16	Chenyi Ang	**Question 1:** We support MAS' proportionate, risk-based approach, including the Annex guidance on assessing whether AI systems are integrated into business processes. To improve consistency across institutions, proportionality may be operationalised through clearer indicators such as reliance on AI outputs, scale of deployment, and potential financial, customer, or regulatory impact of AI errors. Institutions may also find it useful to broadly tier AI use cases (e.g. assistive, decision-support, decision-making) and apply progressively stronger controls as reliance increases. We note that AI systems initially deployed for assistive purposes may become effectively integrated over time due to workflow embedding or scale. Proportionality assessments should therefore be revisited periodically to reflect evolving usage and risk exposure. **Question 2:** **Question 3:** We support the proposed expectations on Board and senior management oversight of AI risk and agree that effective governance requires informed decision-making rather than procedural approvals.

		Boards should understand the purpose of material AI use cases, the level of reliance placed on AI outputs, whether systems exhibit autonomous or agentic behaviour, and the residual risks that remain with the institution. While Boards are not expected to be technical experts, they should be able to challenge management on AI risk and control effectiveness. Existing Board IT or Technology Committees provide a strong foundation but may benefit from explicitly covering AI-specific risks. Institutions may also consider ensuring Boards have access to AI-relevant expertise through training, external briefings, or relevant Board experience. At senior management level, a central AI leadership role (e.g. Chief AI Officer) may support coordination, but does not replace accountability. **Question 4:** We agree that where overall AI risk exposure is material, a cross-functional committee can support coordinated AI governance. In assessing the need for such a committee, institutions may consider AI risk exposure at a portfolio level, including cumulative reliance on AI in critical activities, the presence of Generative or agentic AI, and concentration risk arising from reliance on a small number of vendors. At Board level, institutions may also consider how AI expertise is deployed across existing committees. For example, a Board member with AI or advanced technology experience may sit on multiple committees (e.g. IT, risk, audit) to facilitate coherent oversight and escalation, without creating new Board structures. Such arrangements should be scaled to the institution's size and complexity. **Question 5:** **Question 6:** We strongly support maintaining a comprehensive and up-to-date AI inventory, as this is foundational to effective AI governance. In addition to the attributes outlined in the Guidelines, inventories may usefully capture the level of reliance on AI outputs, whether systems exhibit autonomous or agentic behaviour, deployment models (internal, third-party, embedded), and model update mechanisms. Including monitoring metrics and escalation thresholds further supports lifecycle management, enables senior management and Boards to understand aggregate AI exposure, and facilitates timely response to incidents or supervisory enquiries. **Question 7:** **Question 8:** We support MAS' emphasis on controls across the AI lifecycle, including pre-deployment testing, ongoing monitoring, change management, and appropriate human oversight.

		For higher-risk use cases, such as Generative AI or agentic systems, additional focus may be placed on adversarial testing and red teaming to identify vulnerabilities and unintended behaviours. Monitoring should extend beyond outputs to include actions taken and tools invoked where systems have autonomous capabilities. Clear authority and criteria for suspending or disabling AI systems where risks exceed acceptable thresholds are also important. Ongoing review ensures AI systems remain aligned with their intended purpose and risk appetite as usage and reliance evolve. **Question 9:** AI-related errors or failures should be treated as a form of operational risk. The use of AI, including third-party systems, does not diminish a financial institution's regulatory accountability or duty of care to customers, which remains non-delegable. While contractual indemnities with AI vendors may support commercial risk management, they should be viewed as supplementary and not substitutes for governance and internal controls. Residual risk should be assumed to remain with the institution. Although AI does not constitute a separate risk category, it may materially affect the scale and severity of operational risk. Institutions may therefore consider reflecting AI-related risks within existing operational risk frameworks, including KRIs, scenario analysis, and ICAAP. Where control weaknesses or repeated incidents arise, supervisors may reasonably expect heightened scrutiny and, where appropriate, prudential adjustments under existing frameworks. **Question 10:**
17	Chris Fong Kar Hoe	**Question 1:** • For paragraph 2 of proposed AIRG Annex, specifically 2b., "Is AI integrated with the FI's systems which it is materially dependent on for its business activities?"—the question does not consider if AI integrated in material systems plays a material role or not, or is the intention that as long as any form of AI is integrated into a material system it should be considered as being used as an integrated part of business? More clarity on this would strengthen the guidelines. • For paragraph 5 of proposed AIRG Annex, specifically 5b., basic policies should also include guidelines for: - Addition of AI-specific clauses in all new and existing third party contract / service agreements, at the minimum, clauses for mandatory disclosure of AI usage and obligation to notify significant model update or third party changes, and include other relevant clauses around data ownership, usage and deletion rights, liability and indemnity for AI-related harms, access to audit and transparency reports, mandatory incident notification, data isolation and segregation controls, prohibition of model retraining on client data, as required. This would give FIs the legal basis to gain clarity and manage its third party AI risks as their service providers (i.e. Model provider, AI-enabled SaaS vendor, AI component supplier, Non-AI vendor using AI internally or Non-AI vendor with no AI use) evolve. - Shadow AI management, beyond specifying approved tools and AI use, all FIs should have policy and technical controls (where possible) in place to prevent and detect unsanctioned

		AI use with clear and communicated treatment of such behaviour if detected. - AI literacy, all FIs should provide minimum foundational training on AI technologies, risks and harms to augment their understanding and alignment to basic policies. - Requirements for an established process to assess, prior to each new AI system deployment or extension of deployment, if the deployment alleviates the FI into a position to require more extensive risk management and oversight, and to ensure heightened controls are in place prior to deployment. **Question 2:** • Scope should also consider including AI components, e.g. datasets, embeddings, tools & services (e.g. finetuning / eval tools, prompt orchestration, guardrail services) that feed or influence the FI's AI models and systems. **Question 3:** • On footnote 11 for paragraph 2.3 of proposed AIRG. Consider augmenting the foot note to bring in the context of AI harms when assessing for AI risks. AI harms under the following categories help ground and ensure coverage of responsible AI considerations when identifying AI risks, particularly for the reputational, regulatory, and financial risk categories highlighted in the footnote. Categories of harms include harms to: o Individuals – General harms, e.g. civil rights and safety harm; Specific domain harms, e.g. employment & hiring, insurance & social benefits, housing, education, credit, relationships; Privacy harms, e.g. aggregation, incorporation into training data, inference, secondary use, lack of transparency, inaccuracy; Economic harms, e.g. job displacement; o Groups – e.g. facial recognition bias, mass surveillance; o Society – e.g. mis/disinformation, deepfakes, hallucinations, echo chambers, safety risks, profiling; o Ecosystem – e.g. Natural resources, environmental impact, supply chain distortion; • Additionally, still on footnote 11 for paragraph 2.3 of proposed AIRG, "cultural risks" (AI exceptionalism, over-reliance on AI outputs or automation bias, reduced willingness to challenge AI-generated results, etc.) ought to be highlighted as a risk category as its often-overlooked till it manifests in incidents • On paragraph 2.5 of proposed AIRG, responsibilities of the board (or delegated committee) should include responsibilities to set the overarching vision and strategic direction for AI development, deployment, and usage within the organisation. Ensure AI initiatives align with corporate objectives, risk appetite, and ethical standards. Approves major AI investments and allocates resources for governance and responsible-AI programs, on top of the existing list. This includes ensuring that there is an operational path to translate Responsible AI (RAI) principles subscribed by the organisation (e.g. guided by but not limited to FEAT principles) into system requirements that exhibit Trustworthy AI system (TAI) characteristics, and for baseline AI enterprise risks and AI system harms related to AI development, deployment and usage (in general and specific to the financial sector) to be consistently referenced and analysed for relevance and risk treatment during AI system impact assessments—as part of AI risk management policies and procedures. • Also, on paragraph 2.5 of proposed AIRG, responsibilities of the board (or delegated committee) should include ensuring that effectiveness of the AI governance program itself is measured for progress toward safe value creation and trustworthy AI, for continuous improvement (e.g. Composite AI governance score or maturity index rises year-on-year, %

		reduction in AI incidents and non-compliance cases, % increase in AI initiatives with positive ROI or impact metrics). **Question 4:** • Referencing comments in response to Question 3 on suggested additional responsibilities, that of "Setting the overarching vision and strategic direction for AI development, deployment, and usage within the organisation. Ensuring AI initiatives align with corporate objectives, risk appetite, and ethical standards. Approving major AI investments and allocates resources for governance and responsible-AI programs" should be the responsibility of a cross-functional leadership committee comprised of BSM members, and such a committee should be established for ALL FIs that qualify as "having AI as an integrated part of business processes" • In addition, additional guidelines can be provided on key risk areas where separate cross-functional committees are highly encouraged, with guidance on such committee responsibilities: o AI Oversight Committee – Provides cross-functional oversight of AI governance implementation and ethical alignment. Integrates ethical considerations such as fairness, accountability, transparency, and human-centric values into decision-making and policy review. Monitors compliance with regulations and internal standards, reviews risk assessments and audit findings, and recommends corrective or preventive actions. May draw on internal or external experts for specialised ethical or technical advice as needed. o AI Deployment Review Committee – Evaluates proposed AI systems (internal and third-party) prior to deployment. Assesses their ethical, technical, and operational suitability, monitors post-deployment performance, and recommends continuation, modification, or decommissioning of AI systems based on governance outcomes. o AI Vendor Risk Assessment Committee – Evaluates and monitors third-party AI vendors for cybersecurity, data-protection, and ethical-risk considerations. Maintains an approved-vendor list and conducts periodic reassessments based on performance and compliance outcomes. **Question 5:** • Systems, policies and procedures to ensure the consistent identification of AI usage across all relevant business and functional areas does not necessarily ensure interpretation to include prevention, detection and actions for detected shadow AI or unsanctioned AI use—particularly when guided concurrently by proposed AIRG paragraph 4.3 establishing AI Inventory as a control for unapproved usage of AI. Proposed AIRG can augment clarification that AI identification includes both AI use cases and AI tools use arising from official avenues (which gets updated into AI inventory) as well as unsanctioned ones (i.e. Shadow AI), in foot notes under proposed AIRG Section 3 or directly in wordings of paragraphs 3.2 and 3.3 of proposed AIRG. • Ideally, in alignment with basic policy recommendation under response for Question 1, a separate guideline paragraph should be dedicated to Shadow AI management under proposed AIRG section 3—beyond specifying approved tools and AI use, all FIs should have policy and technical controls (where possible) in place to prevent and detect unsanctioned AI use with clear and communicated treatment of such behaviour if detected. **Question 6:**

		• On paragraph 3.5 of proposed AIRG, AI inventory should not only capture key attributes of the AI model, system or tool, it is also critically, a repository where key AI risk management process artefacts (and other key software development process documentation) are kept for each AI model, system or tool (e.g. project context documentation, impact assessment, Trustworthy AI characteristics translation to system requirements, model/system card documentation, metric selection and TEVV tooling documentation, testing and evaluation results, third-party AI vendor due diligence documentation etc.). This will allow the AI inventory to enable effective governance and oversight, as well as risk management AND auditability **Question 7:** • The current dimensions—impact, complexity and reliance—can be very open for interpretation and seem overly focused on just model risk rather than AI systems as a whole. A set of dimensions, tied to RAI first principles and grounded in real-world system behaviour and consequences would have greater applicability and provide greater clarity needed for AI system risk classification when organisational use cases do not fit neatly into the lists highlighted in the guidelines. This set of dimensions should come with guidelines (in Annex) on materiality levels which then form the basis for a quantitative scoring system that can be applied consistently across FIs, regardless of use-case. The resulting risk score category—low to critical—then guides the structured approach for proportionate application of lifecycle controls. Risk Materiality Dimensions: o Potential Harm Severity - Level 1, Minimal: Cosmetic or fully reversible issues with no impact on fundamental rights or safety. Represents the lowest risk threshold. - Level 2, Limited: Reversible inconvenience or low financial impact. May cause temporary disruption but no lasting consequences. - Level 3, Serious: Material harm or rights infringement becomes possible. Requires enhanced monitoring and intervention capabilities. - Level 4, Critical: Irreversible injury, significant rights violations, or systemic organizational loss. Demands highest level of oversight. o Autonomy & Reversibility - Level 1, Manual Assist: Fully human-initiated actions with complete reversibility. AI provides recommendations but humans make all decisions. - Level 2, Semi-Automated: Human-in-the-loop for key decisions. System can act but requires human approval for critical actions. - Level 3, Autonomous: Operates without real-time oversight. Decisions are made independently with periodic human review. - Level 4, Critical: Cannot be undone once triggered. Represents highest risk level requiring strongest safeguards. o Data Sensitivity & Rights Impact - Level 1, Non-Sensitive: Public or synthetic data with minimal privacy implications. - Level 2, Internal: Confidential business data but not personally identifiable. - Level 3, Personal: Contains identifiable user or customer information. - Level 4, Special Category: Contains identifiable user or customer information. o Operational Criticality - Level 1, Low: Non-critical business process. - Level 2, Medium: Impacts performance, not safety.

		- Level 3, High: Affects safety, compliance, or continuity. - Level 4, Critical: Mission or life-critical infrastructure. o Scale & Reach - Level 1, Narrow: Limited to few users within isolated deployment context. Pilot projects or small team tools typically fall here. - Level 2, Moderate: Departmental or pilot use across multiple teams. Affects dozens to hundreds of users within controlled boundaries. - Level 3, Broad: Enterprise-wide or national exposure. Thousands of users or decisions affected across organizational or geographic scope. - Level 4, Massive: Global or cross-sector systemic influence. Millions affected with potential cascading effects across industries or societies. o Robustness & Cybersecurity Posture - Level 1, Strong: Verified security controls, defence-in-depth architecture, regular penetration testing, and comprehensive monitoring. Robust incident response procedures in place. - Level 2, Moderate: Partial controls implemented with some known vulnerabilities. Basic monitoring and testing conducted periodically but gaps exist. - Level 3, Weak: Reactive security posture with limited proactive testing. Vulnerabilities may be unaddressed and monitoring capabilities are minimal. - Level 4, Fragile: No resilience mechanisms or monitoring in place. System is highly vulnerable to attacks, manipulation, or failure with no detection capability. o Human Oversight Strength - Level 1, Full Oversight: Human reviews and approves all key outcomes before execution. Maximum control and accountability maintained. - Level 2, Partial Oversight: Human reviews high-risk cases or statistically sampled decisions. Automated decisions proceed with spot-checking. - Level 3, Limited Oversight: Human notified of decisions but cannot block execution. Monitoring is informational, not preventive. - Level 4, No Oversight: Fully automated, self-executing decisions with no human involvement. Highest risk level requiring strictest controls. Total Risk Score Tier Classification: o Tier 1 – Low (7-11): Minimal harm potential; low autonomy; non-sensitive data. o Tier 2 – Moderate (12-16): Some sensitive data or moderate autonomy/scale. o Tier 3 – High (17-22): Significant rights/safety impact, high autonomy or critical ops. o Tier 4 – Very High / Critical (23-28): Safety-critical or systemic risk. **Question 8:** • On paragraph 4.9 of proposed AIRG, it would provide more clarity if under footnote 19 or in an additional footnote, example guidance can be provided on what protected attributes may constitute (e.g. e.g. Race/ethnicity, Political opinions, Religious/philosophical beliefs, Trade union membership, Genetic/biometric data, Health data, Sexuality, sexual orientation) • On paragraph 4.11 On Third-party AI management, consider adding the following under key areas that FI should consider: o AI-specific vendor classification – FIs should develop and maintain an AI-specific vendor classification framework that integrates with existing vendor management processes where possible. The framework could categorise vendors as AI vendors (e.g. AI Model Provider, AI-Enabled SaaS Vendor, AI Component Supplier), Non-AI Vendors Using AI

		Internally, and Non-AI Vendors with No AI Use, to distinguish their respective roles within the AI supply chain. The classification enables proportionate application of due diligence, contractual obligations, and AI risk management controls according to the nature and level of AI involvement. o Vendor Due Diligence – FIs should conduct AI-specific due diligence on all vendors according to their classification to ensure appropriate assessment of AI-related risks. Due diligence shall cover relevant domains including Transparency & Documentation, Data Governance, Bias & Fairness, Security & Resilience, Privacy & Data Protection, Explainability, Accountability & Oversight, Compliance & Certification, Third-Party Dependencies, Ethical & Legal Commitments, Incident Management, and End-of-Life & Exit Strategy. The scope and depth of review shall be proportionate to the vendor's AI classification and risk exposure. o Alignment with Risk Management Policy - For AI Models and AI Components, FIs should apply full lifecycle risk management activities on systems which the models or components are integrated with. For AI-Enabled SaaS, only partial application of the AI Risk Management Policy may be feasible as these may be complete systems with limited configurability or access to underlying models and data. For such cases, FIs may rely on third-party accreditation under international AI risk management system standards, e.g. ISO/IEC 42001 and outcomes of vendor due diligence for areas pertaining to testing, evaluation, verification and validation (TEVV), bias and fairness testing, training-data governance, and knowledge-limit documentation • On paragraph 4.16, Technology and Cybersecurity Risks, besides relevant regulatory requirements and guidance relating to technology risk management, FIs should also take into account relevant industry frameworks/guidelines, e.g. the OWASPs top 10 security threat mitigation guidelines for ML, LLM, Agentic applications, and MCP, and continuously update their organisational AI security mandates as new industry guidelines are published. Also, in alignment with the culmination of the latest OWASP guidelines, the following can be considered in replacement of the current areas of security, access control and third-party as a more holistic set of AI security principle pillars for FIs to follow: o Enforce a Zero-Trust, Least-Privilege Architecture – Assume no implicit trust; continuously verify every component and enforce minimal permissions by default to constrain the blast radius of any compromise. o Ensure Model, Data and Third-party AI Component Integrity Across the Lifecycle – An AI system's reliability is dictated by the integrity of its data and components. The entire AI supply chain (internal or external) must be secured particularly around models (internal, third-party closed / open weights models), and AI components that feed or influence the AI systems (libraries, plugins, and datasets etc., including third-party / open-sourced ones) o Implement Continuous Monitoring and Human-in-the-Loop Oversight – AI systems demand non-negotiable visibility and control. Implement comprehensive monitoring to understand AI system behaviour and mandate human intervention for all high-stakes decisions. o Build for Resilience and Containment – Assume preventative controls will fail. Systems must be architected with safeguards to contain incidents, prevent runaway processes, and ensure operational stability. • On paragraph 4.17, Reproducibility & Auditability, focus on documentation should go beyond replicating AI system or model and its results which addresses mainly model risks. It is also critically important for enabling and evidencing risk management and oversight, i.e. providing traceability and evidence of the rationale and key decisions made along the

		AI lifecycle. The proposed AIRG can consider adding additional paragraphs for the following: o Standardise AI System Documentation – Adopt a standardised template, such as a "Model Card" or "Fact Sheet," for every AI system. This ensures that critical information is captured consistently and is accessible to all relevant stakeholders (e.g. revision history, model name & version, point of contact, business justification, intended uses & scope, training data description, description of output data, algorithmic methodology—model architecture and training approach, evaluated alternative approaches—other models and methods considered and rationale for choice, down and up-stream dependencies, testing & validation results, known limitations & assumptions, stakeholder engagement plans, deployment & monitoring plans). o Document and maintain AI risk management artefacts that provides context and demonstrates key risk management decisions (e.g. Project context worksheet, AI system risk classification, Impact assessment, Trustworthy AI system requirement mapping, Metric and TEVV tooling selection documentation, Testing & validation results and follow-up, Model card/Factsheet, Stakeholder engagement matrix, Deployment & monitoring plans), uploading or referencing them as part of the AI system record in the organisation's AI inventory. • On paragraph 4.23, particularly 4.23b., Incident and Issue Management, consider adding guidance on AI-specific incident classification. FIs should develop and maintain an AI-specific incident classification framework that integrates with existing incident management processes where possible. An AI-specific incident classification framework provides categories that better describe the unique nature of AI incidents allowing for more effective incident identification, classification, articulation of root causes, and establishment of guidelines for incident response actions and response times. AI incident classification may include: o Cybersecurity incidents—model or data compromise, adversarial manipulation, supply-chain threats; o Unauthorized or unexpected outcomes—from hallucination, goal drift, emergent behaviour, poor prompt or tool design in agentic systems, weak boundary controls; o Discrimination/fairness harms; o Privacy violations; o Physical/health & safety incidents; o Transparency/accountability failures—from model opacity, poor documentation or audit trail; o Performance degradation/model decay. **Question 9:** The following are highlighted as aspects of AI risk management that are not covered in the proposed AIRG, in addition to those touched on in responses to earlier questions. AI Risk Management Systems, Policies & Procedures Layer Component: • AI Literacy – AI literacy should be a key component under Key AI Risk Management Systems, Policies & Procedures. Requiring training on AI foundational understanding, responsible AI principles and trustworthy AI system characteristics, AI risks and harms, organizational AI policies—particularly on approved AI tools, AI usage and repercussions of shadow AI, proposal of AI use case and feedback loops. Such training should be provided to all hands as basis for further, more specific BSM or role-based training for necessary AI capabilities and capacity building (which guidelines cover well at each area)

		• In addition to existing paragraphs 1.8 and 1.9 in proposed AIRG which covers guidance on enterprise AI risks which FIs should take into account, consider adding additional paragraph(s) for guidance on AI system harms and expand the "Risk" section to cover "Risk and Harms". AI harms in the following categories ground and ensure coverage of responsible AI considerations when identifying AI risks, particularly for the reputational, regulatory, and financial risk categories listed in 1.8 and 1.9. AI harms may include harms to: o Individuals – General harms, e.g. civil rights and safety harm; Specific domain harms, e.g. employment & hiring, insurance & social benefits, housing, education, credit, relationships; Privacy harms, e.g. aggregation, incorporation into training data, inference, secondary use, lack of transparency, inaccuracy; Economic harms, e.g. job displacement; o Groups – e.g. facial recognition bias, mass surveillance; o Society – e.g. mis/disinformation, deepfakes, hallucinations, echo chambers, safety risks, profiling; o Ecosystem – e.g. Natural resources, environmental impact, supply chain distortion; • "Cultural risks" (i.e. AI exceptionalism, over-reliance on AI outputs or automation bias, reduced willingness to challenge AI-generated results, etc.) ought to also be highlighted as a risk category under paragraph 1.9 as an additional key risk from Generative AI being an often-overlooked risk till it manifests in incidents. • Baseline Enterprise AI risks and AI system harms registry – In addition to proposed AIRG guidance on "AI identification", "AI inventory" and "AI Risk Materiality" as key components of an FI's AI risk management system (AI RMS), organisations maintaining a registry of baseline AI enterprise risks and AI system harms related to AI development, deployment and usage is also a key AI RMS component and should be included under Section 3 of the guidance. A baseline AI risk and harms register serves as the foundation for FIs to perform gap analysis on their existing policy framework and determine what additional standalone AI-related policies and existing AI-specific extensions of existing policies are needed. It also serves as a basis for consistent assessment of risk and harm relevance when conducting necessary AI impact assessments for each AI project undertaken by FIs. • Policy and procedural development – While policies and procedures are mentioned frequently in the proposed AIRG, it is primarily done in the context of the topic the guidance is addressing. Having an AI risk management policy framework and architecture is a key component of an AI RMS and guidance should be included, under Section 3 of the proposed AIRG, on the need for an established process together with guidance on key areas for policy coverage: o Conduct a policy gap assessment to map existing policies against identified AI enterprise risks and determine overlaps or gaps. o Policy coverage strategy guidance: - Stand-alone AI policies: For areas introducing new risk dimensions or regulatory obligations (e.g., AI Governance & Oversight, Responsible AI & Ethics, AI Risk Management, AI Security, AI Procurement, Intellectual Property & Data Ownership, AI Use, and AI Training & Competence). - AI-specific extensions: For existing policies that can be expanded to include AI elements (e.g., Data Governance & Privacy, Cybersecurity, HR & Recruitment, Vendor Risk Management, Intellectual Property & Information Management, Audit & Compliance). AI Life Cycle Controls Layer Component: • Context Establishment – Document AI system purpose, business value, deployment setting, potential beneficiaries, foreseeable risks and harms (referencing organisation

		baseline AI risk and harms registry), and any applicable regulations the system may be subject to. This first step of the AI lifecycle / risk management process sets the foundation for later risk management activities and guides FIs to uncover information that flows into the AI impact assessment and other risk management artefacts e.g. model cards / factsheets. • AI Impact Assessment (not to be co-mingled with AI risk materiality assessment) – AI system risk classification is assessed based on dimensions related to some nature or characteristic of the AI model, system or tool, and is meant to guide proportionate application of controls and oversight. While the AI impact assessment takes in project context and system requirements for identification of foreseeable risks and harms, evaluates and documents their potential severity, assesses for risk treatment (i.e. avoid, reduce, transfer, accept) against organisational risk appetite, and evaluates and documents residual risks post mitigation approaches, including how residual risks are monitored post deployment. Paragraphs 3.8 through 3.11 describes AI risk materiality assessments as an AI system risk classification control, and there is no further guidance in the guidelines on the need or AI impact assessments to be conducted for AI projects. Both AI system risk classification and AI Impact Assessment are key AI lifecycle / risk management controls, assessing AI systems with different lenses, and the AI system risk classification control itself is insufficient for holistic identification of AI project risks for evaluation, treatment, and identification residual risk monitoring. • Translate Trustworthiness into System Requirements – The Responsible AI (RAI) principles an organisation adopts (e.g. FEAT but not limited to) should drive the organisation's adoption of relevant Trustworthy AI (TAI) characteristics it wants to see in the AI systems it develops, deploys and use (e.g. Privacy-enhanced, Explainable & interpretable, Fair & harmful bias mitigated, Secure & resilient, Safe, Valid & Reliable, Accountable & transparent—as prescribed by NIST AI RMF, and Regulatory Compliant—where applicable). To ensure AI systems exhibits these TAI characteristics, an established process is needed to elicit system requirements that map to them, e.g. a TAI system requirement mapping questionnaire that guides the product owner or solutions architect in evaluating if a TAI characteristic is relevant to the AI project and what system requirements is needed. • Usage of Open-Source Components - Ensure usage of open-source models and datasets are reviewed and approved in alignment with organisational policy prior to integration into AI project. Policies on usage of open-source components should include: compatibility of licensing terms with organisational and commercial objectives; data provenance, privacy, and consent obligation; and documentation and enforcement of obligations for attribution, redistribution, or use restrictions. • Stakeholder Engagement – AI systems affect people beyond direct users, risk emerges across the lifecycle, and no single function—legal, risk, data, security, or engineering—has full visibility or ownership of AI risk on its own. FIs should develop a stakeholder engagement framework that clearly identifies relevant stakeholder categories (e.g. Strategic governance stakeholders; Assurance, risk, legal, compliance, and security functions; Technical builders and operators; Business owners and internal users; External vendors and AI suppliers; and External impacted parties), specifies their role (i.e. informed, consulted, leads, approves, or some trigger-specific involvement beyond these four core roles), and establishes the lifecycle triggers for stakeholder communication (e.g. New AI use case proposed, AI risk classification, Data/model/vendor decisions, TEVV model validation, Pre-deployment review & approval, Deployment & rollout, Operational

		monitoring & drift, Incident response, Periodic recertification, Regulatory reporting, etc.). This ensures that stakeholder engagement is activated by lifecycle triggers and not left to personal judgment. **Question 10:**
18	Clearspring Insights	**Question 1:** We concur with the intent of proportionate application and welcome the examples of non-integrated use cases provided in the Annex. However, we note that the definition of "integrated part of business processes" in Paragraph 2 appears to rely heavily on operational dependency (i.e., disruption of workflow) rather than risk materiality. For example, an internal IT helpdesk chatbot may be operationally useful, but its failure presents low risk compared to a credit scoring or payment fraud detection model. We respectfully suggest that the Authority consider generalising the exemptions into a broader "Assistive Tool Principle". We propose that AI tools be excluded from the "integrated" definition if they meet the following criteria: - Non-Customer Facing: The tool does not interact directly with external customers; - Human-in-the-Loop: A human operator retains full discretion and decision-making authority over the output; and - Reversible & Low Impact: Errors are reversible by the human operator and do not result in material financial impact or regulatory breaches. This distinction ensures the framework remains robust for high-risk systems (e.g., algorithmic trading or real-time payment sanctions screening) while appropriately exempting productivity enhancers. **Question 2:** We are of the view that the current scope might inadvertently stifle early-stage experimentation, particularly for GenAI and AI Agents. To align with MAS' strategic priority of "broadening adoption" and "anchoring leading AI capabilities," it is vital to distinguish between live production systems and internal R&D. We encourage MAS to formally recognise "Internal Innovation Safe Harbours". AI models used solely in non-production, sandbox, or internal R&D environments should be explicitly exempt from inventory, validation, and monitoring requirements until they are promoted to production or interact with live customer data. This would allow FIs to iterate rapidly on new technologies without premature compliance burdens. **Question 3:** We seek clarification on the requirement for the Board to have an "adequate understanding of AI" (Para 2.5d). We respectfully submit that this should be interpreted as an understanding of Strategic Risk and Impact Tolerance, rather than technical engineering proficiency.

		To assist Boards in meeting this requirement, we recommend the Guidelines explicitly allow Boards to rely on "Competency Augmentation"—leveraging external advisors or specialised internal committees to bridge technical knowledge gaps. This aligns with the MAS Guidelines on Risk Management Practices (Board and Senior Management), which permit Boards to delegate detailed risk oversight to specialised committees or advisors while retaining ultimate accountability. **Question 4:** We recognise MAS' assessment that AI, when material, represents a unique and significant risk vector that may warrant dedicated oversight parallel to established pillars such as Audit or Risk. However, unlike Audit (which is a distinct line of defense), AI is a transverse enabler embedded within business lines (e.g., Credit, Fraud, Operations). We respectfully submit that mandating a separate "AI Committee" risks treating AI risk in a vacuum, divorced from the business context where the actual impact occurs. Therefore, we propose that FIs be permitted to leverage "Integrated Risk Management (IRM)" as an alternative to a separate committee. Under this model, FIs would mandate a Dedicated AI Agenda within existing Board Risk Committees, supported by the cross-functional expertise described in Question 3. This ensures that AI risk is given the specific weight MAS desires, but remains tightly coupled with the holistic risk management of the firm. **Question 5:** We note that requiring "robust systems" to identify all AI usage (including what the industry terms "Shadow AI" or the "unapproved usage" cited in Para 4.3, such as web-based GenAI tools) implies a need for invasive surveillance that may conflict with operational agility. We propose a "Materiality-First" identification standard: - High-Impact AI: Must be robustly identified and tracked via automated scanning. - Low-Risk/Shadow AI: FIs should be permitted to use "Risk-Based Sampling"—a methodology endorsed in existing MAS AML/CFT circulars—to monitor for unauthorised AI use, rather than mandating 100% real-time coverage. This places the onus on FIs to prove their sampling methodology is effective without imposing disproportionate surveillance costs. **Question 6:** We feel that a "comprehensive" inventory for every minor AI tool may create administrative noise that obscures true risk. We propose a Tiered Inventory Model: - Tier 1 (High Risk/Material): Detailed tracking (lineage, validation status, owner) for AI that affects customer outcomes or has material financial impact.

		- Tier 2 (Assistive/Low Risk): Aggregate reporting (e.g., “Generic Productivity Tools”) with periodic reviews. We also suggest that inventory updates be “Change-Triggered” (e.g., upon promotion to production or material model change) rather than “real-time,” which is more practical for operational teams to maintain. **Question 7:** We agree with the dimensions of Impact, Complexity, and Reliance and recommend two additions to capture modern systemic risks: Systemic Concentration: Explicitly assessing reliance on a small oligopoly of third-party foundational model providers (Hyperscalers), which poses a systemic resilience risk similar to cloud concentration risk in payments. Reversibility: Acknowledging that “reversible” AI errors (e.g., a draft email reviewed by a human) carry significantly lower risk than “irreversible” ones (e.g., an automated payment execution). This distinction is critical for proportionate risk scoring. **Question 8:** Regarding Explainability (Para 4.6): We highlight that for GenAI, “intrinsic explainability” is often technically impossible. We urge MAS to explicitly accept “Outcome-Based Evaluation” (Red Teaming, functional testing) and “Traceability” (citations/RAG) as valid substitutes for explainability. Regarding Third-Party Audit (Para 4.11): Securing “right to audit” clauses with global hyperscalers is commercially unrealistic for individual FIs, creating a compliance deadlock. We respectfully request MAS to formally recognise “Third-Party Certifications” as acceptable mechanisms to satisfy oversight requirements. Specifically, we recommend the Guidelines reference the following standards established by international bodies as valid proxies for direct audits: - ISO/IEC 42001 (Artificial Intelligence Management System): Developed by the International Organization for Standardization (ISO), this is the first international standard specifying requirements for an AIMS. Crucially, ISO 42001 operationalises the risk management protocols of ISO/IEC 23894 and impact assessments of ISO/IEC 42005 into a certifiable framework. It effectively serves as an operational "container" for these controls, providing FIs with verified assurance of a vendor's governance maturity. - SOC 2 Type 2: Developed by the American Institute of Certified Public Accountants (AICPA), a Type 2 report evaluates the operational effectiveness of controls over a period (typically 6-12 months). This provides FIs with empirical validation that a vendor's claims regarding security, availability, and processing integrity are accurate, without requiring a direct audit of the vendor's infrastructure. **Question 9:**

		We recommend refining the guidance on "kill switches" (Para 4.23). Drawing from financial market controls, we suggest distinguishing between appropriate controls for Autonomous Agents versus GenAI content tools: - Kill Switches: Essential for autonomous agents or algorithmic trading systems where immediate cessation is necessary to prevent financial loss. - Circuit Breakers/Graceful Degradation: For GenAI information tools, the system should detect anomalies and divert traffic to human agents or rule-based fallbacks rather than terminating the service entirely. This ensures business continuity while mitigating risk. We also suggest adopting a Shared Responsibility Model (similar to the framework standardised by the Cloud Security Alliance (CSA) for cloud computing) to clearly delineate liability between Model Providers (core safety) and FIs (application context). **Question 10:** Given the complexity of retrofitting governance onto legacy systems and the nascence of GenAI controls, we respectfully submit that a mandatory 12-month deadline for full compliance is impractical for the industry at large. We support a Phased Transition of 24 Months to ensure quality implementation: - Phase 1 (12 Months): Framework setup and remediation of High-Materiality systems. - Phase 2 (24 Months): Full remediation of legacy systems and Low-Materiality tools.
19	Commerzbank Aktiengesellschaft Singapore Branch	**Question 1:** **Question 2:** **Question 3:** **Question 4:** **Question 5:** **Question 6:** **Question 7:** **Question 8:** **Question 9:** **Question 10:** The bank kindly requests that MAS consider allowing FIs to use EU AI Act implementation and internal risk management programs as evidence of capability building. Additionally, MAS may consider a structured, risk-based implementation roadmap (with milestones) during the transition period, rather than requiring full implementation across all AI use-cases within 12 months.

20	CPA Australia	**Question 1:** Singapore's financial sector already benefits from strong technology foundations and high-quality regulatory frameworks. CPA Australia therefore does not consider significant regulatory changes necessary. Any regulatory changes should remain both proportional and risk-based. High-risk AI systems such as those involved in lending decisions, insurance underwriting or fraud detection should attract more stringent requirements, while lower-risk applications such as customer service chatbots or internal productivity tools should be subject to lighter oversight. Including SME considerations is essential. Small businesses often lack the resources to implement complex governance and assurance frameworks, and imposing uniform requirements risks creating excessive compliance cost without materially reducing systemic risk. A proportionate approach ensures obligations are calibrated to the size, risk profile and operational complexity of each organisation, supporting financial inclusion while enabling SMEs to innovate responsibly. Regional alignment also strengthens proportionality. ASEAN's Responsible AI Roadmap (2025–2030) and updated ASEAN Guide on AI Governance and Ethics advocate a risk-based, principles-driven approach. Aligning MAS's approach with these frameworks will enhance interoperability, reduce cross-border compliance fragmentation and reinforce Singapore's leadership in responsible AI. International harmonisation remains important, but global regimes are still evolving. Singapore should pursue interoperability pragmatically. **Question 2:** High-risk AI systems such as those used in lending decisions, insurance underwriting and fraud detection should receive the strongest oversight. Lower-risk use cases — such as customer service chatbots or internal productivity tools — should attract lighter requirements. CPA Australia's Business Technology Survey 2025 found that while 92 per cent of businesses in Singapore report using AI, most use it on an irregular basis and for productivity support rather than automating critical financial operations. This suggests that MAS should adopt a clear but proportionate scope that distinguishes between high-risk, high-impact AI and lower-risk applications. **Question 3:** Trust in AI depends on robust governance and clear accountability within financial institutions. Senior management and boards must understand AI's role, risks and ethical implications, embedding AI oversight within existing risk frameworks rather than treating it as purely technical. Institutions should maintain comprehensive documentation of training data, model design and decision logic to support supervisory review and consumer transparency. Human accountability remains critical. High-impact decisions must retain meaningful human oversight to ensure technology-enabled processes do not erode consumer

		protections. As highlighted in CPA Australia's submission to IOSCO, responsibility for AI-driven processes must sit with people, not systems. **Question 4:** CPA Australia supports MAS's proposal for cross-functional governance committees where AI risk is material. Such committees ensure that AI oversight is not isolated within technical teams but is shared across business, compliance and operational functions. CPA Australia suggests clarifying indicators or thresholds for what constitutes "material" AI risk. Clear, standardised triggers — especially for high-risk AI use cases — would help institutions determine when such a committee is required and ensure consistent application across the industry. **Question 5:** Institutions should maintain comprehensive documentation of training data, model design and decision logic. This supports supervisory review and transparency, and it naturally contributes to the consistent identification of AI use cases across business and functional areas. While your submission does not identify specific definitions, it emphasises the importance of robust systems, lifecycle controls and clear governance structures — all of which rely on being able to identify when and where AI is deployed. **Question 6:** Institutions should maintain comprehensive documentation of training data, model design and decision logic. Maintaining such documentation supports the development of an accurate and up-to-date inventory of AI systems across the organisation. This enables consistent oversight, accountability, and risk management throughout the AI lifecycle. **Question 7:** CPA Australia supports a risk-based regulatory approach. High-risk AI systems such as those used in lending decisions, underwriting or fraud detection clearly align with impact, complexity and reliance dimensions. These systems present significant consumer and operational risks and should therefore attract stronger requirements. Lower-risk applications such as chatbots have lower impact and complexity, and should receive lighter oversight. Obligations must be calibrated to the size, risk profile and operational complexity of each organisation, particularly SMEs. This is consistent with ASEAN's Responsible AI Roadmap and Singapore's risk-based regulatory philosophy. **Question 8:** AI systems should undergo rigorous testing before deployment. Ongoing monitoring and validation are equally important to assess accuracy, fairness, resilience and security, especially given that AI models evolve with new data. Independent assurance strengthens confidence in AI systems. High-risk AI systems should include kill switches and manual fallback mechanisms.

		MAS can effectively regulate AI by enhancing existing frameworks rather than creating standalone regimes: Technology Risk Management Guidelines: can incorporate AI-specific controls Outsourcing Guidelines: should include obligations for provider transparency, bias testing and fallback mechanisms Cyber Hygiene Notices: should expand to include AI vulnerabilities such as adversarial attacks FEAT Principles: embedding FEAT into supervisory processes and conducting explicit FEAT compliance checks would strengthen fairness, transparency and explainability Enhancing existing frameworks minimises duplication and enables proportionate, principles-based regulation. **Question 9:** The effectiveness of AI regulation depends on the regulator's ability to keep pace with technological developments. MAS should invest in targeted training, recruitment and tools such as regulatory sandboxes and model validation frameworks. Capacity building must extend to financial institutions, especially SMEs, through templates, workshops and phased implementation milestones. MAS could align its guidelines with national AI initiatives. Existing programmes such as IMDA's AI adoption grants, Enterprise Singapore's productivity schemes and the Enterprise Compute Initiative provide financial support and technical training for businesses. Coordinating with these efforts would reduce duplication and reinforce Singapore's whole-of-government approach to responsible AI. Public confidence is essential. Institutions should clearly disclose when and how AI influences decisions, explain the factors involved, and provide clear pathways for human review and appeal. Dispute resolution processes should be accessible and well-communicated. Financial literacy programs should expand to include AI literacy. MAS could also consider minimum transparency standards for high-risk AI, including plain-language explanations of AI's role, disclosure of key decision factors, and clear instructions for contesting decisions. **Question 10:** CPA Australia supports MAS's proposed 12-month transition period. This should be supplemented by phased milestones to support SMEs, who may need additional time and capability-building to implement governance, assurance and lifecycle controls proportionately and effectively.
21	Curia Regis Pte Ltd	**Question 1:** The Annex's examples of AI use in business processes are helpful. However, we suggest adding fund management-specific examples such as AI-enabled portfolio optimisation and rebalancing tools used systematically, AI models for risk factor analysis integrated into

		investment processes, AI-driven ESG scoring systems for portfolio screening, Robo-advisory platforms providing discretionary portfolio management, etc. While we understand the current wording allows for this, but this can be quite generic whereas it would be appropriate to identify the areas where more prescriptive guidance may be useful. **Question 2:** First, regarding third-party data and analytics providers: Fund managers commonly use third-party platforms (e.g., Bloomberg, Refinitiv, MSCI) that increasingly incorporate AI/ML features. We request clarification on whether such embedded AI features fall within scope where the fund manager does not have visibility into the underlying AI models. This relates to existing outsourcing requirements under the MAS Guidelines on Outsourcing (MAS-G07). Second, regarding quantitative investment strategies: Traditional systematic and quantitative strategies may use statistical models that border on machine learning. We suggest MAS provide clearer delineation between traditional quantitative methods and AI/ML for the purposes of these Guidelines, potentially referencing the definition approach used in MAS' existing Model Risk Management guidance for banks. Third, regarding AI used by service providers: Fund administrators, custodians, and transfer agents increasingly deploy AI in their operations. We request clarity on fund managers' obligations regarding AI used by their service providers, particularly given existing obligations under the FMC Guidelines (paragraph 5.6) to ensure service providers are competent and capable. **Question 3:** We recommend MAS clarify that the oversight responsibilities may be fulfilled through documented processes and periodic reviews rather than requiring formal Board committee structures. **Question 4:** The fund management industry in Singapore includes many smaller operations where a dedicated cross-functional committee may be disproportionate. Unlike banks with multiple business lines and large operational functions, a typical fund manager may have a lean structure comprising portfolio management, compliance, and operations functions only. This may be more relevant for FMCs that are above S$1 b in AUM. Additionally, for the assessment of overall AI risk exposure, we suggest MAS consider the following fund management-specific factors: Whether AI directly influences investment decisions for discretionary mandates, extent of AI use in suitability assessments for advisory services, reliance on AI for compliance monitoring (e.g., investment restrictions, regulatory limits), use of AI in client-facing communications, AI involvement in trade execution and best execution obligations **Question 5:** We recommend MAS permit fund managers to leverage existing operational risk identification and technology asset registers rather than requiring entirely separate AI-specific systems. This aligns with the proportionate application principle and recognises that many fund managers already maintain comprehensive IT asset inventories under their

		Business Continuity Planning requirements (FMC Guidelines paragraph 5.4) and Technology Risk Management obligations (MAS-TRM Guidelines). The designated control function responsible for AI identification should, where appropriate, be the existing Compliance function which already has responsibilities for monitoring adherence to regulatory requirements and internal policies. **Question 6:** We support maintaining an AI inventory and recommend that the inventory requirements be integrated with existing registers that fund managers are already required to maintain: 1. Outsourcing Register: Under the MAS Guidelines on Outsourcing, fund managers maintain records of material outsourcing arrangements. AI provided by third parties could be integrated into this existing framework. 2. Technology Asset Register: Under TRM Guidelines, fund managers maintain inventories of technology systems. AI systems naturally fit within this existing framework. **Question 7:** **Question 8:** **Question 9:** **Question 10:** We recommend MAS consider a phased transition period: - 12 months for basic policies and governance frameworks, and AI identification and inventory completion, - 24 months for full AI life cycle controls implementation
22	Cyber Risk Institute	**Question 1:** CRI supports MAS's principle of applying these Guidelines proportionately across financial institutions (FIs), as outlined in paragraph 1.5 and the accompanying Annex. This approach acknowledges the diverse landscape of AI adoption and risk profiles within the financial sector. MAS's guidance rightly differentiates between institutions requiring "basic policies" for limited AI use, and those needing more comprehensive frameworks, policies, and procedures when AI is an "integrated part of their business processes." This aligns with CRI's core belief that AI risk management should be structured based on a firm's level of AI adoption. Both approaches recognize that risk management expectations should scale with the complexity, criticality, and interconnectedness of AI development, deployment, and use. CRI's AI Adoption Stage Questionnaire may serve as a complementary tool to MAS's current Guidelines on proportionality. Our Questionnaire was designed to guide FIs in a systematic self-assessment of their current AI adoption stage. It classifies firms into one of four stages, providing a more granular and objective basis for proportionate application than a binary classification. The 4 stages are: • Initial: Characterized by limited and protective AI methods. • Minimal: Involves low-risk implementations. • Evolving: Entails high-risk production applications. • Embedded: Describes AI widely integrated across the organization.

		These stages consider critical dimensions such as Business Impact, Technology Implementation, and Scalability (among others). By evaluating statements describing current practices across these dimensions, FIs can objectively determine their AI Adoption Stage. MAS's distinction between 'basic policies' and 'integrated use' provides a valuable foundational approach for AI governance. For financial institutions, CRI's structured, multi-stage adoption framework, as detailed within the FS AI RMF, can serve as a complementary tool for internal calibration. This four-stage model offers additional granularity, enabling FIs to more precisely identify their specific AI adoption position and thus calibrate their risk management efforts with greater accuracy and proportionality, while remaining aligned with MAS's foundational guidance. The AI Adoption Stage Questionnaire also provides a standardized, repeatable mechanism for FIs to assess their current and future stage of AI adoption. This structured self-assessment helps objectively determine the appropriate scope and depth of their risk management systems, policies, and procedures, thereby operationalizing MAS's proportionality guidance. The CRI stages naturally link to risk materiality (e.g., "minimal" for low risk, "evolving" for high-risk production), directly supporting MAS's expectation that FIs use AI as an integrated part of business processes to "apply clear identification and robust risk materiality assessments." By adopting a robust, multi-stage AI adoption framework that combines MAS's Guidelines and CRI's Adoption Stage model, organizations can gain a clearer understanding of their current state, better apply the MAS Guidelines proportionately, and more objectively prepare for future AI use. This approach would not only help FIs meet regulatory expectations but also support their responsible innovation journey with greater transparency and effectiveness. **Question 2:** CRI's approach in developing the FS AI RMF was to remain use case agnostic, deliberately casting a wide net to provide a broad foundation for AI risk management and governance for financial institutions, particularly those in the earlier stages of AI adoption. However, we see significant value in MAS's proposed scope of AI use cases, systems, and models to apply the Guidelines. We commend MAS for its comprehensive approach, recognizing its robust coverage of risk categorization, ethical considerations, monitoring and evaluation frameworks, data quality standards, staff training and capabilities, incident response, and the provision of specific AI use case examples. These elements collectively form a strong foundation for responsible AI adoption in the financial sector. To further enhance these Guidelines, we offer a few additional thoughts for consideration. Drawing from our observations of emerging best practices, we suggest that while the Consultation Paper outlines stakeholder feedback mechanisms for deployed AI, FIs might benefit from proactive stakeholder engagement during their own AI system development lifecycle. Additionally, as FIs increasingly operate across borders, providing clearer guidance on cross-jurisdictional AI regulatory considerations would be valuable. Beyond leveraging parent entity frameworks, insights into navigating diverse and potentially

		conflicting AI regulatory landscapes in different international markets could support global compliance efforts. **Question 3:** We appreciate MAS's emphasis on the pivotal role of the Board and senior management in governing AI-related risks, fostering an appropriate risk culture, and ensuring robust frameworks are in place. There is strong alignment between MAS's proposed expectations—such as approving the overall governance approach, integrating AI risks into the risk appetite framework, and ensuring adequate capabilities and resources—and CRI's thinking on defining clear roles and responsibilities, establishing AI risk appetite and tolerance thresholds, and fostering a risk-aware culture. MAS's call for regular reviews to keep pace with AI developments and for clear communication channels is very important and well-placed. To further embed accountability, FIs themselves could consider integrating clearer AI risk management responsibilities into individual job functions. While the Guidelines broadly outline Board and senior management duties, defining how AI risk oversight translates into specific roles (e.g., AI system owners, developers, data scientists, risk managers, independent validation teams) could foster deeper internal accountability. MAS might consider emphasizing the importance of FIs clearly delineating such internal responsibilities. The Guidelines rightly point to the need for "appropriate training and capacity building." FIs could enhance their capabilities by implementing role-specific AI risk management training programs for relevant personnel. Such training could focus on trustworthiness principles, integrating content lineage and provenance, and be complemented by advanced certification programs for specialized roles. MAS might consider encouraging FIs to develop evaluation mechanisms to assess training effectiveness and ensure practical application of knowledge. To minimize "risk blind spots" and groupthink, best practices suggest that FIs strive to assemble AI risk management teams that deliberately prioritize multidisciplinary perspectives—considering different demographics, disciplines, experience, and backgrounds. While recognizing that the feasibility of fully diverse teams may vary by FI size and resources, encouraging a broader, more inclusive approach to identifying and managing AI risks could enhance robustness. Lastly, while human oversight is mentioned in the Guidelines, more direction on defining and managing "Human-AI Supervision Policies" could be valuable. This would address risks like automation complacency, anthropomorphization, and disinformation, establishing safeguards for human override and promoting transparency in human-AI interactions. For more examples related to the above, consider reviewing CRI's FS AI RMF risk and control matrix, specifically the control objectives listed in the Govern function, Categories 2, 3, 4, and 5. **Question 4:**

		We support MAS's proposal for establishing a dedicated cross-functional committee for AI risk oversight when an FI's exposure is material. This is consistent with CRI's view on forming robust AI governance structures with senior leaders from various functions to provide oversight. MAS's approach to assessing inherent and residual risk materiality, and linking it to the FI's risk appetite, is a sound foundation. While a cross-functional committee is an excellent starting point, MAS might consider suggesting the evolution of this governance structure towards clearer, singular accountability, potentially involving a dedicated senior executive who holds primary responsibility for AI risk. This would ensure alignment with existing risk functions and elevate AI risk management to a strategic C-suite level topic. For more examples related to the above, consider reviewing CRI's FS AI RMF risk and control matrix, specifically the control objectives listed in the Govern function, Categories 2, 3, 4, and 5. **Question 5:** We appreciate MAS's clear proposal for financial institutions to establish robust processes for consistently identifying AI use cases, systems, and models. This foundational step is essential for effective governance and risk management throughout the AI lifecycle. We recognize that overly prescriptive definitions of AI can quickly become outdated, given the technology's rapid evolution. The Guidance's emphasis on a principled approach, such as the use of FEAT, for defining and characterizing AI is well-aligned with this reality. While the Guidance appropriately acknowledges the evolving nature of the technology (e.g., in Section 2.9), CRI's FS AI RMF offers a complementary layer of operational detail. We believe there is an opportunity to collaborate on how the principled approach within the Guidance maps to more granular frameworks, like the FS AI RMF, to offer financial institutions a common language for both high-level principles and their flexible, practical application. This collaborative mapping could further enhance consistency where possible, while ensuring definitions remain relevant over time. While "robust systems" are mentioned, the Guidelines might offer more specific recommendations on technological enablers for AI identification. This might include suggestions for automated scanning tools, advanced AI inventory platforms, or tools to differentiate AI components and dependencies within complex systems. CRI's risk and control matrix describes multiple control objectives in Govern 1.6. **Question 6:** MAS covers essential attributes like purpose, data, dependencies, lifecycle status, and risk materiality, aligning well with CRI's objectives for robust AI inventory management. We are very closely aligned with respect to core inventory elements and lifecycle management, integration with broader "enterprise" risk management, third-party AI, and risk-based resourcing for inventory. Building on the Guidelines' mention of decommissioning, FIs could enhance their AI inventory requirements by explicitly including specific decommissioning attributes such as

		sunset dates, archival procedures, and data retention policies. This would ensure more thorough end-of-lifecycle management. Additionally, specifying the integration of the AI inventory with existing technology management and data governance processes could drive broader consistency and cross-functional collaboration. Further, encouraging the analysis of inventory data for common risk patterns, interdependencies, and cascading effects across AI systems. could enable more strategic, enterprise-wide risk mitigation rather than just individual system-level management. It might also be beneficial to provide clear guidance for identifying and incorporating AI in "Shadow IT" and purchased products. Such guidance could help FIs prioritize robust IT/Architecture review of all purchases to ensure business units are not acquiring AI solutions without proper oversight. Lastly, proposing regular, independent audits of the AI inventory itself is key. Advocating for automated tools to support inventory processes could significantly enhance these efforts. Such audits would verify completeness, accuracy, and timeliness, ensuring reliability as a consistent source of truth for AI governance. Embedding a process to integrate "lessons learned" from operational monitoring, audit, and other feedback directly into inventory management processes might also facilitate continuous improvement and adaptation of the inventory framework over time (CRI's risk and control matrix describes multiple control objectives in Govern 1.6). **Question 7:** These dimensions, particularly impact and complexity, align with industry standards and the approach taken in CRI's own maturity model, which emphasizes the need for organizations to determine how these foundational dimensions apply within their specific operational context. **Question 8:** The Guidelines—seeking comments on proposed standards, processes, and controls across the entire AI lifecycle, and the call for proportionate application based on relevance—are comprehensive, timely, and relevant. This broad view is critical, as practical challenges often arise in coordinating roles, responsibilities, and hand-offs between various functional units throughout the AI lifecycle. Establishing clear governance structures and processes is paramount to address these operational complexities effectively. While the Guidance provides clear expectations, financial institutions may require more specific resources and tools to design and implement appropriate AI risk programs. MindShare has already developed a framework for this purpose and is currently creating operationalization resources. We recommend CRI's FS AI RMF as an additional, complementary resource. Developed through extensive public-private collaboration with more than 100 financial institutions, and aligned with the NIST AI RMF, the FS AI RMF offers detailed, non-prescriptive control objectives and practical implementation guidance that can assist firms in applying these guidelines proportionately across their evolving AI journeys. We would welcome collaborating on aligning and mapping these initiatives to provide a comprehensive suite of tools for the industry. **Question 9:**

		The Consultation Paper does a good job addressing third-party AI risks within the Guidelines, particularly in Section 4.11, which outlines key considerations such as transparency, fairness, supply chain assessments, concentration risks, contingency plans, and legal agreements for third-party AI. This foundation aligns well with many of our control objectives related to third-party AI evaluation, contracting, and general monitoring. To further strengthen this critical area, we propose a few enhancements. Consider: 1. Explicitly addressing fourth-party risks. For critical third-party AI, additional guidance could emphasize AI-specific risk transfer mechanisms, such as AI insurance policies or contractual provisions, within legal agreements and contingency planning. 2. Further emphasis on AI-specific training for personnel managing third-party relationships, focusing on the unique nuances of AI due diligence, contract negotiation, and ongoing risk assessment. 3. Formalized audits for third-party compliance, going beyond general monitoring, to enable verification of adherence to AI-specific data management, security standards, and contractual obligations. Overall, we encourage MAS to review CRI's FS AI RMF and use it as a resource benefiting both financial institutions and regulators. FIs can leverage its structured, comprehensive control objectives and scalable implementation guidance tailored to their AI adoption stage, identify specific risks, and integrate AI risk management into existing enterprise risk and governance frameworks. For regulators, the FS AI RMF offers a common, industry-backed taxonomy and detailed set of control objectives, providing a robust benchmark for evaluating FIs' AI risk management practices, informing supervisory expectations, and fostering consistent evaluation across the financial sector. This shared understanding can drive more effective risk mitigation while promoting responsible AI innovation. Please see CRI's risk and control matrix for additional insights, specifically, Govern 6.1 and 6.2, Map 4.1 and 4.2, and Manage 3.1. **Question 10:**
23	Cyber Sierra (Fort One Technologies Pte. Ltd.)	**Question 1:** 1. Clarify Thresholds for "Integrated" AI Use: The Guidelines should more clearly define when an FI is considered to be using AI as an "integrated part of business processes." While paragraph 1.5 states that FIs must adopt AI policies proportionate to their level of AI use and the Annex distinguishes ad-hoc from integrated usage, clearer criteria or thresholds (e.g. reliance of critical systems or workflows on AI) would reduce ambiguity, especially for smaller FIs. This approach aligns with international practices: the U.S. Federal Reserve applies simpler model risk frameworks where model use is limited, and the EU AI Act adopts a risk-based model that imposes the strictest requirements only on high-risk AI systems. 2. Introduce Tiered Compliance Expectations:

		The Guidelines could adopt a tiered structure of requirements, differentiating baseline controls versus advanced controls based on the FI's size or AI risk exposure. Currently, MAS proposes that "oversight and AI risk management… may be proportionate to the size and nature of the FI's activities", but providing concrete tiers (for example, categorizing FIs or AI use cases into low, medium, high risk tiers) would guide implementation. This approach would align with international best practice. The EU AI Act, for example, explicitly categorizes AI systems by risk level (minimal, limited, high, unacceptable) and calibrates obligations accordingly. By defining tiers (perhaps in the Annex), MAS can ensure that smaller FIs or low-risk AI use cases implement basic governance (as outlined in Annex paragraph 5), while large FIs or high-impact AI deployments implement comprehensive controls. 3. Provide Tools for Small FIs: MAS could consider providing templates or checklists to assist smaller FIs in implementing the "basic policies" required under paragraph 1.5. These basic policies (e.g. designating a responsible senior manager and rules on acceptable AI use) are outlined in the Annex. To ensure consistent adoption, MAS might publish a guideline supplement or toolkit for smaller FIs, covering minimum documentation, oversight structures, and training for staff. This practical assistance would echo approaches in other jurisdictions – for instance, Canada's OSFI has issued detailed guidance (Guideline E-23) for all banks on model risk management, including AI, with expectations scaled by the institution's size and complexity. By helping smaller firms with ready-made frameworks, MAS can achieve proportionate compliance in practice, not just in principle **Question 2:** 1. Refine the Definition of "AI" in Scope: The Guidelines may benefit from clearer definition of what constitutes an AI use case, system or model, to avoid confusion at the margins. Currently, MAS defines AI broadly to include machine learning (ML), deep learning, reinforcement learning, Generative AI, agents, etc., while explicitly excluding tools based solely on fixed logic or rules. This definition could be refined by clarifying, for example, that adaptive or data-driven decision-making is the hallmark of "AI" in scope. It would also help to align with global definitions. The EU AI Act defines an AI system as software developed with "machine learning, logic- and knowledge-based approaches, or statistical approaches" that can generate outputs for given objectives. Notably, the EU includes advanced expert systems under AI, whereas MAS excludes pure rule-based software. MAS should confirm if exclusion of rule-based systems is intentional (to focus on learning systems), and if so, articulate this rationale in the Guidelines. This clarity will help FIs uniformly identify which technologies are governed by the AI risk framework. 2. Include Third-Party and Vendor-provided AI: The scope should explicitly cover AI systems acquired or used from third-party providers, not just in-house models. Paragraph 1.2 already signals that the Guidelines apply across different AI technologies and will be updated for new developments. MAS could strengthen this by stating that any AI models or services that an FI uses (whether built in-

		house or obtained via vendors or cloud APIs) fall under the Guidelines' purview. This ensures comprehensive coverage of AI risk, given many FIs leverage external AI solutions. It also aligns with global expectations on third-party risk – for example, the FSB has highlighted "over-reliance on a few dominant AI providers" as a key emerging risk. By explicitly including third-party AI, MAS reinforces that FIs must exercise the same rigor in oversight (due diligence, testing, monitoring) for vendor AI as for internally developed models. 3. Clarify Treatment of Generative AI and AI Agents: The Guidelines mention that newer AI paradigms like Generative AI and AI agents are in scope, but additional guidance on their use cases would be useful. MAS should clarify that even experimental or pilot uses of Generative AI (e.g. using large language models for internal tasks or customer chatbots) are subject to appropriate risk assessment. The consultation's Preface notes challenges like hallucinations, IP issues, and privacy leaks from Generative AI – it would be prudent for MAS to explicitly require FIs to factor these into their risk controls if they deploy such AI. Similarly, AI agents (software with autonomy to take actions) might pose new types of operational risk. Detailing in the scope that any AI with autonomous decision-making authority, including agents that can act on behalf of the FI, must be governed by these Guidelines will help FIs include these cutting-edge tools in their inventories and oversight mechanisms. This forward-looking stance ensures the scope remains relevant as AI technology evolves. **Question 3:** 1. Enhance Board Accountability and Expertise: The Guidelines should emphasize that Boards are ultimately accountable for AI risk and must develop sufficient understanding of AI to exercise effective oversight. Section 2.5 already assigns the Board (or a delegated committee) responsibilities such as approving the AI risk governance framework and ensuring AI risks are within the FI's risk appetite. An improvement is to require Boards to receive specialized training on AI risks and ethics, or to include members with relevant expertise. This aligns with IOSCO's guidance that firms should designate senior management with the "relevant skillset and knowledge" to oversee AI, with proportionality based on the AI's impact. It also mirrors the BIS's finding that governance and skills are areas needing further regulatory attention in AI risk management. By explicitly stating that Board members should attain adequate AI knowledge (as is hinted in 2.5(d)), MAS would reinforce effective challenge and informed oversight at the top. 2. Foster a Culture of AI Risk Awareness: MAS could strengthen the point that Board and senior management must cultivate the right risk culture regarding AI use. Paragraph 4.1 already notes they should "foster the appropriate risk culture for the use of AI". MAS might consider to expand this by suggesting that tone-from-the-top on ethical AI and risk vigilance is essential. For example, the Board should set expectations that AI use should align with the FI's core values and FEAT principles (Fairness, Ethics, Accountability, Transparency). Senior management should operationalize this through training programs for staff on AI risks and ethical use, and by

		incorporating AI risk considerations into performance and incentive structures. This is in line with the OECD AI Principles, which stress accountability and human-centric values in AI governance (principles that MAS has adopted via FEAT). By embedding AI risk awareness into the corporate culture (e.g. regular internal communications from leadership on responsible AI, encouraging staff to speak up about AI concerns), FIs can more effectively manage AI risks in day-to-day operations, beyond formal policies. **Question 4:** 1. Define the Mandate and Composition of the AI Risk Committee: The Guidelines should offer more detail on the expected role and makeup of the dedicated cross-functional AI risk committee for FIs with material AI exposure. MAS proposes establishing such a committee when overall AI risk is material, but FIs would benefit from clarity on its functions. An improvement is to specify that this committee should coordinate AI governance across business units, review significant AI initiatives, monitor aggregate AI risk, and report to senior management/Board regularly. Its composition should draw from multiple disciplines – e.g. risk management, IT, compliance, business lines, data science, and audit – to ensure all perspectives are represented. This echoes emerging best practices: many leading banks have formed AI governance committees that include legal, compliance, model risk, IT security, and business heads. NIST's AI framework also recommends "cross-functional governance" for AI, often through dedicated committees. By outlining mandate and membership (perhaps in an Annex or new paragraph), MAS will help FIs set up effective committees rather than mere formalities. 2. Set Criteria for "Overall AI Risk Exposure" Assessment: To determine when the AI risk committee is needed, MAS should provide a structured approach for assessing overall AI risk at the enterprise level. Footnote 15 of the consultation suggests considering the number of high-risk use cases, deployment in critical business lines, and impact on the FI's risk profile. MAS could expand this into a multi-factor scoring or qualitative criteria. For example, FIs might evaluate: (a) Impact – potential financial, reputational, customer harm from AI failures (borrowing from the risk materiality dimensions in paragraph 3.10); (b) Pervasiveness – proportion of processes or decisions powered by AI; (c) Complexity/Novelty – use of opaque models like deep learning in core areas; and (d) Third-Party Dependence – reliance on external AI providers. If thresholds in these areas are exceeded, overall AI risk is "material." Providing such guidance aligns with practices in model risk management: U.S. regulators encourage an enterprise model risk aggregate view and often require banks to have a "Model Risk Committee" once model usage reaches a certain scale. Likewise, a clear threshold (e.g. if any high-risk AI is in a critical service, or X% of credit decisions use AI) would prompt formation of the committee. This ensures consistency across FIs in determining when heightened oversight is warranted. 3. Transparency and External Input: To strengthen trust and accountability, MAS might encourage FIs to allow some form of independent input or review into the AI risk committee's activities. For instance, the committee could invite Internal Audit or risk experts not directly involved in AI

		development to provide challenges (paragraph 4.18 already suggests independent pre-deployment reviews for high-risk AI). Additionally, MAS could suggest that where appropriate, the committee consults external guidance – e.g. ethical AI advisors or industry consortium findings – especially on emerging risks. This resonates with global moves: the EU AI Act will require appointing a compliance officer for AI in high-risk cases, and firms are advised to involve diverse stakeholders in AI governance. While a formal external member on the committee may not be feasible for all, MAS could recommend periodic external audits or peer reviews of the committee's effectiveness. This external perspective will help the committee stay objective and aligned with industry best practices. At a minimum, MAS could require FIs to document the committee's remit and activities clearly and be prepared to share these with MAS supervisors, ensuring regulatory visibility into how overall AI risk is being overseen at the top levels. **Question 5:** 1. Adopt a Formal Definition of "AI Use Case/System" Internally: The Guidelines should urge FIs to create an internal definition or taxonomy of what constitutes an "AI use case, system or model" for identification purposes. Paragraph 3.2 already expects FIs to have "clear definitions, criteria and processes" for consistent AI identification. As an improvement, MAS could suggest that FIs document a formal definition (aligned with MAS's scope) that employees can easily apply. For example, an FI might define AI systems as "any IT system or process that employs machine learning or automated decision logic which adapts based on data." This should be communicated via policy so that business units know what to flag as AI. The U.S. Fed's model risk guidance took a similar approach by first defining what counts as a "model" across the firm (including qualitative approaches that produce quantitative outputs), thereby ensuring all relevant tools are captured in inventories. By explicitly requiring a definition and enterprise-wide understanding, MAS will help avoid situations where pockets of AI usage go unrecognized because they weren't labeled as such by the business. 2. Use Attestation and Certification Processes:To ensure completeness, MAS could suggest that FIs implement attestation processes where business heads regularly certify that usage of AI in their domain has been identified and logged. Paragraph 3.3 hints at this by assigning a control function to set up "attestation processes" and act as final arbiter of what is AI. Reinforcing this, MAS might require an annual (or semi-annual) sign-off: each business and support unit declares AI-driven tools or confirms none are used. This top-down check compels management to think through their operations and consult the definition of AI. It mirrors practices in areas like model risk and EUC (end-user computing) tools governance, where business units attest to their inventories. Such attestations, overseen by the designated control function, would bolster consistency. The control function (perhaps residing in Technology or Risk Management) can challenge and validate these attestations – e.g. cross-checking with procurement records, project lists, or even keyword scans of code bases for AI libraries. This process ensures that identification isn't only passive but actively verified. It aligns with MAS's expectation that a control function "act as the final arbiter" in determining AI usage, thereby enforcing consistency across the FI. **Question 6:**

		1. Specify Minimum Attributes for the AI Inventory: The Guidelines should enumerate the key information fields that each AI inventory entry should contain, to ensure completeness and consistency. Paragraph 3.5 lists examples of attributes (purpose, scope, model type, data, dependencies, lifecycle status, risk rating, validation status, key roles, etc.). MAS could elevate this by making some of these attributes mandatory for all FIs. For instance, every AI asset in the inventory should record: a unique identifier, a description of its use, the business owner, the risk materiality classification, last validation date, and whether it's third-party or in-house. Setting a baseline schema would help smaller FIs know where to start and allow MAS supervisors to compare inventories across FIs more easily. This aligns with practices in model risk management – the OCC in the US expects that "a firm-wide inventory" of models is maintained with comprehensive information (purpose, ownership, status, etc.). By defining the core fields (perhaps in an Annex or as guidance), MAS ensures each inventory is not just a list of names, but a rich repository supporting governance and oversight as intended. 2. Ensure the Inventory Covers Development and Decommissioned AI: MAS should clarify that the inventory must include AI systems under development and those recently retired, not only those currently deployed. This is hinted at in best practices (e.g., the US Fed notes inventories should include models "under development… or recently retired"). Including AI in development helps risk managers and validators engage early (before deployment), and tracking retired AI ensures lessons learned (and potential re-use of components) are documented. The Guidelines could be explicit: FIs should update the inventory at key milestones – when an AI project enters development, when it is approved for deployment, and when it is decommissioned. Paragraph 3.4 does say new, updated or decommissioned AI should be reflected accurately; MAS can reinforce this by expecting inventory entries to have status labels (e.g. "Under Development", "In Production", "Decommissioned") and last updated timestamps. Maintaining historical records of AI usage is important for audit and model risk capital assessments. It also prepares FIs for potential future regulations (for example, the EU AI Act will require providers to maintain technical documentation and logs even after an AI system is discontinued). Emphasizing full lifecycle coverage in the inventory will lead to a truly "accurate and up-to-date" record as MAS envisages. 3. Integrate Inventory Maintenance into Change Management: To keep the AI inventory current, MAS should instruct FIs to embed inventory updates into their change management and project governance processes. In practice, this means whenever a new AI system is deployed or an existing one is significantly changed or retired, the change approval process should require an inventory update as a checklist item. Section 4.25 of the Guidelines covers managing changes to deployed AI (defining significant vs minor changes, etc.). MAS could add that as part of change management, the inventory must be reviewed and amended. For example, if an AI model is retrained with a major methodology update (a "significant change"), its inventory entry should be updated with the new model version, date, and any change in risk assessment. Likewise, if a use case is expanded to a new jurisdiction or customer segment, that scope change is reflected

		in the inventory attributes. Embedding these steps formalizes inventory upkeep. It also resonates with global expectations that inventories be living documents. The FRB's model risk guidance implicitly assumes continuous inventory updates by requiring banks to have an inventory and robust model change control processes. By tying inventory updates to the existing change management controls in the Guidelines, MAS ensures the inventory doesn't become stale and accurately mirrors the live AI environment. **Question 7:** 1. Include "Data Quality/Sensitivity" as a Dimension: In addition to impact, complexity, and reliance, MAS should consider adding a risk dimension explicitly addressing data quality and sensitivity. Paragraph 3.10(a) touches on data in the Impact dimension (noting nature and sensitivity of data processed), but elevating data to its own category could sharpen focus on risks like personal data misuse or poor training data quality. This Data dimension would consider factors such as: the sensitivity of input data (e.g. personal identifiable information, which raises privacy and compliance risk), the representativeness and quality of data (which affects bias and accuracy), and the potential for data drift over time. These aspects are critical in AI outcomes – for example, biased or unrepresentative data can lead to unfair decisions (a risk to consumers and reputational risk to the FI). Other regulators emphasize data issues: the EU AI Act imposes strict data governance requirements for high-risk AI, including that datasets be "relevant, representative, free of errors and complete". IOSCO's AI guidance likewise highlights "data quality and bias" as a key risk area to assess separately. By explicitly making data a risk dimension (or splitting the impact dimension into Impact – Customers/Financial vs Impact – Data), MAS ensures FIs give due weight to data-related vulnerabilities in their materiality assessment. 2. Add "Third-Party/External Dependency" as a Dimension: While the current Reliance dimension partly covers reliance on AI autonomy and alternatives, it would be beneficial to explicitly capture third-party risk. Many AI systems rely on external providers (cloud AI platforms, pretrained models, open-source libraries). A Third-Party Dependency dimension would consider concentration risk (e.g. the FI's dependence on a single vendor for critical AI) and the risk from vendor practices (security, reliability, and ethics of third-party models). MAS's consultation acknowledges this in principle – e.g. new risks like "over-reliance on a few dominant Generative AI providers" are noted but including it in the formal assessment framework will ensure FIs evaluate it systematically. For instance, two AI models might be equally complex and impactful, but one built in-house vs. one entirely API-driven from a tech giant have different risk profiles due to external dependency. Other regimes also focus on this: the FSB flagged "AI-related third-party dependencies and concentration" as a potential systemic risk factor. Thus, incorporating a third-party dimension (or explicitly mentioning it under Reliance) would prompt FIs to assess supply-chain risks of AI – such as what happens if a vendor's service fails, or if an open-source model has unknown vulnerabilities. 3. Incorporate an "Ethics/Fairness" Dimension:

		MAS could introduce an additional dimension addressing ethical considerations and fairness of AI outcomes. The current Impact dimension does mention fairness and ethical breaches as considerations, but given the importance of these issues (also reflected in MAS's FEAT principles), making it a standalone dimension could be justified. This Fairness & Ethics dimension would prompt FIs to assess risks like discrimination, lack of transparency, or misalignment with societal values. For example, an AI lending model might be low-complexity and internally developed (low on complexity and reliance), but if it poses high fairness risk (potentially biased against certain groups), that should elevate its materiality. Including this dimension aligns with global best practice: fairness and transparency are core tenets of trustworthy AI frameworks (OECD, EU, etc.), and the EU AI Act's notion of risk explicitly ties to impact on fundamental rights (which covers discrimination and human oversight). By isolating ethical risk, MAS reinforces that technical efficacy alone is not the sole concern, social impact matters too. In practice, some FIs already conduct fairness impact assessments for AI; making it a formal part of risk materiality scoring would integrate those efforts into overall risk management. .4. Clarify Weighting and Interaction of Dimensions: In improving the framework, MAS could provide guidance on how FIs should combine the dimensions of impact, complexity, and reliance (and any new ones) to derive an overall risk materiality rating. Right now, 3.10 says the assessment should "consider various risk dimensions… minimally covering" those three, but FIs might struggle with how to weigh them. MAS might suggest an approach: for example, consider Impact as the primary factor (since it captures potential harm), then use Complexity and Reliance as modifiers that can raise the risk level (since higher complexity or autonomy can exacerbate the impact risk). Alternatively, MAS could endorse a straightforward high/medium/low scoring on each dimension and then define what combinations imply high overall materiality (akin to worst-case driver). This guidance would promote consistency across FIs. It would also reflect how other frameworks treat multi-dimensional risk: the NIST AI RMF, for instance, encourages assessing likelihood and severity of harms (analogous to impact), but also considering controllability and context (similar to reliance/complexity). By clarifying how to synthesize the dimensions, possibly through examples or a matrix in an Annex, MAS can help FIs perform more rigorous and uniform risk materiality assessments, and ensure that if any additional dimensions (like fairness or third-party) are added, they are integrated coherently into the overall risk rating process. **Question 8:** 1. Strengthen Requirements for Robust Testing (including Stress Testing): The Guidelines cover evaluation and testing of AI (section 4.14) and suggest that testing be proportionate to risk. An improvement is to explicitly require stress testing of AI models, especially high-risk ones, under a range of adverse scenarios. For example, FIs should test how AI systems perform under data drift, extreme market conditions, or when inputs are intentionally perturbed (to simulate adversarial attacks or unusual but plausible scenarios). This is hinted in MAS's footnote for Generative AI, which advises attention to issues like hallucinations and adversarial attacks in evaluation. Making stress testing a standard control (where relevant) ensures AI models are not just validated on historical data but assessed for resilience. IOSCO's guidance similarly urges firms to "carry out adequate testing… in stressed and unstressed conditions.". By incorporating stress testing into the

		life cycle controls (perhaps by expanding 4.14 or in a new clause), MAS will push FIs to anticipate and mitigate failure modes of AI, thus preventing issues like procyclical behaviors or breakdowns during crises. 2. Mandate Transparency to Users and Affected Customers: The Guidelines emphasize internal transparency and explainability (sections 4.6, 4.7), but MAS could also encourage external transparency, i.e. informing customers or users when AI is making significant decisions about them. Currently, there's no explicit requirement to disclose AI usage to customers in the Guidelines, yet this is increasingly seen as best practice for consumer protection. For example, if an AI declines a loan or advises an investment, customers might be informed that an algorithm was involved and given an explanation of key factors. This concept is supported internationally: the EU AI Act will require user disclosures for certain AI (like chatbots or emotion recognition) and the EU's GDPR already gives individuals rights around automated decisions. IOSCO likewise suggests regulators consider requiring firms to "disclose pertinent information to customers and clients around their use of AI/ML", especially where it affects client outcomes. MAS could incorporate a proportional approach – e.g. for high-impact customer-facing AI (credit, insurance underwriting, robo-advice), FIs should implement procedures to explain decisions and notify customers that AI is used, in line with existing laws (like PDPA's advisory guidelines on AI use). Including this in the life cycle controls (perhaps under Transparency & Explainability or Fairness) would strengthen accountability and customer trust without being overly prescriptive, since lower-risk internal AI might not need such disclosure. 3. Embed Ethical and Human-Centric Controls Across the Life Cycle: The Guidelines individually mention fairness (4.8–4.9), human oversight (4.10), and explainability (4.6–4.7) – all of which are excellent. MAS could tie these together under an overarching principle that AI systems should be "human-centric, fair, and accountable" throughout their life cycle. For instance, during development (Section 4.12), developers should document how they addressed fairness and bias; during evaluation (4.15), include fairness metrics and explainability tests; during deployment, ensure human review for high-stakes AI (4.10); and during monitoring, track outcomes for any unfair impacts (4.23(e) suggests user feedback mechanisms). By framing these not just as discrete controls but as a continuous mindset, MAS aligns the life cycle controls with the FEAT principles that underpinned earlier MAS work. Internationally, frameworks like the OECD AI Principles and NIST AI RMF emphasize a cross-cutting need for fairness, transparency, and accountability at each stage of AI's life. Concretely, MAS could require FIs to conduct an Ethical AI assessment before deployment (somewhat covered by fairness assessment in 4.9) and to have documented justification that the AI's benefits outweigh potential ethical risks (especially for novel AI). Additionally, MAS could highlight the importance of involving diverse stakeholders or committees (e.g. an ethics committee or customer advocacy input) when developing AI for customer-facing purposes, aligning with the trend of multidisciplinary AI governance. Infusing this human-centric perspective across the AI life cycle controls will ensure technology serves customers and society responsibly, not just efficiently. **Question 9:**

		1. Address Regulatory Compliance and Legal Liability of AI: One aspect not explicitly covered is how FIs should manage compliance risk arising from AI, and clarify liability when AI decisions breach laws. The Guidelines focus on risk management processes, but MAS could explicitly remind FIs that adoption of AI does not absolve them from meeting all existing legal and regulatory requirements (as noted in 2.4: Board must ensure compliance with relevant regulations even when AI is adopted). An improvement is to provide guidance that FIs should assess how AI systems comply with laws on consumer protection, anti-discrimination, data privacy (PDPA), etc., during their development and deployment. For example, if an AI is used in credit scoring, the FI should ensure it does not result in unfair lending in violation of fair lending laws. MAS might also clarify that in the event an AI causes a regulatory breach or customer harm, the FI will be held accountable just as if a human made the decision. This reinforces the "accountability" principle, the onus is on FIs to control their AI. Such clarity aligns with approaches elsewhere: US regulators (e.g. CFPB) have warned that using AI or algorithms is not a defense for non-compliance with consumer protection laws, and the UK's FCA has mapped how its existing rules (e.g. treating customers fairly) apply to AI use. By integrating compliance checks into the AI governance framework (perhaps requiring a legal/compliance sign-off in the AI development lifecycle), MAS ensures FIs proactively manage legal risks and understand that "the buck stops with them" for AI's actions. 2. Strengthen the Role of Internal Audit and Independent Review: While the Guidelines mention independent pre-deployment reviews for high-risk AI (4.18–4.19), they do not explicitly discuss the ongoing role of Internal Audit or other independent assurance functions post-deployment. MAS could add that Internal Audit should periodically audit the AI risk management framework and selected AI use cases for adherence to the Guidelines. This provides a third line of defense to identify gaps. For instance, Internal Audit could assess whether model validations are effective, whether monitoring controls are operating, or whether the AI inventory is complete (as discussed earlier). MAS itself may rely on such internal audit reports during supervision. A comparison: The Federal Reserve's model risk guidance explicitly expects Internal Audit to "assess the overall effectiveness of the model risk management framework". Likewise, MAS could expect audit to cover AI governance. Additionally, FIs could be encouraged to leverage external experts for independent model validation or ethical audits of AI (especially for novel AI uses) to supplement internal capabilities. This was not directly covered in the Guidelines but could be included as a good practice. Having robust independent review mechanisms will provide greater assurance that AI controls are truly working as intended and not just on paper. 3. Consider AI Risk Disclosure to Regulators or Public: An area not covered is transparency of AI use at a systemic level. MAS could consider whether FIs should disclose certain information about their AI adoption either to MAS regularly or even publicly. For example, MAS might collect data through inspections or surveys on the material AI use cases in key sectors (banking, insurance) to inform its supervisory strategy. Over time, MAS could even contemplate requiring FIs to include a section in annual reports or Pillar 3 disclosures on their AI governance (similar to how

		banks disclose operational risk or cyber risk preparedness). This is a forward-looking suggestion. Currently, few jurisdictions mandate AI disclosures, but the EU AI Act will create a public EU database of high-risk AI systems (albeit focused on providers). Some global banks voluntarily report on their responsible AI initiatives in sustainability reports. Pushing for greater transparency could enhance stakeholder trust. In formal feedback, we propose MAS at least begin with supervisory disclosures: e.g., FIs might report any "significant AI incidents" to MAS (paralleling tech outage notifications under TRM), and annually inform MAS of the top AI risks they see. Such information sharing wasn't explicitly mentioned in the Guidelines but would greatly assist MAS in monitoring industry AI risk trends and calibrating its oversight (as recommended by the FSB for authorities to close data gaps in AI monitoring). **Question 10:** 1. Emphasize Proportionality During Transition: MAS could clarify that during the transition period, it expects FIs to make good-faith, risk-based efforts toward compliance, focusing first on the highest risk AI activities. In practice, this means an FI with one critical AI model and several trivial ones should prioritize bringing the critical model under the new governance within the first few months. MAS might communicate (via a supervisory circular or FAQ) that it will adopt a consultative supervisory stance in the first year, helping FIs work through challenges rather than immediately penalizing deficiencies, provided there is a credible implementation plan. This approach echoes the implementation of other major frameworks (for example, when MAS rolled out past guidelines like TRM, smaller FIs were guided to prioritize key controls). Proportional transition expectations will alleviate industry anxiety and aid innovative initiatives within the FI. The outcome should be that by the 12-month mark, all FIs have the foundational elements in place (governance structure, inventory, initial risk assessments), and are well on their way to full adherence. If certain FIs need slightly more time for specific complex controls, MAS could handle that via case-by-case extensions. Signaling this flexibility, without formally changing the 12-month target, can ensure FIs focus on risk outcomes (like no unchecked high-risk AI running) rather than mere compliance formality by the deadline. 2. Provide Implementation Guidance or Templates Early in the Period: To help FIs meet the 12-month transition, MAS could publish additional guidance, templates, or examples soon after finalizing the Guidelines. For instance, a template for an AI risk inventory schema, or a sample risk materiality assessment framework, or a mock charter for an AI risk committee. These would not be one-size-fits-all, but illustrative aids that FIs (particularly those with limited resources) can adapt. MAS could also share anonymized, aggregate best practices observed from early adopters or pilot firms. This proactive support will reduce trial-and-error time. It's akin to how some regulators issue implementation FAQs or organize industry workshops right after a major rule, effectively catalyzing the compliance process. Given the tight timeline, if MAS finalizes the Guidelines (say by Q1 2026), hosting a roundtable with industry within a couple of months to address questions will be valuable. Another idea is to encourage FIs to perform a gap analysis as the first step (comparing current AI governance against each section of the Guidelines) – MAS could provide a checklist for this. By scaffolding the transition with guidance, MAS

		increases the chances that all FIs, not just the large ones, hit the ground running and achieve compliance by the 12-month mark.
24	CyberArk Software Pte Ltd	**Question 1:** Alignment with Identity Security: This proportionate approach makes sense and aligns with identity security best practices. Baseline identity controls should be in place for any AI use, regardless of institution size. For example, at minimum, every FI that deploys AI should classify AI systems as a form of digital identity and enforce least-privilege access for them. MAS's guidance that "all FIs should minimally institute basic policies for the use of AI commensurate with [their] level of AI adoption" supports this by ensuring even small FIs implement essential controls (like defining who oversees AI use and what uses are allowed). Larger or more AI-intensive FIs would scale up to full lifecycle risk management, similar to how they scale identity governance programs. Gaps/Recommendations: To strengthen this principle: Baseline Controls for All: MAS could explicitly enumerate the minimum baseline controls that every FI must implement, even in a lightweight manner. For instance, every FI should at least maintain an inventory of AI applications and apply least-privilege access controls to AI systems, no matter how limited their AI use. This guarantees a floor of security. Examples in Annex: The Annex could provide examples of proportionate implementation. For example, a small FI with one or two AI-based tools might have a single accountable officer and a simple register of those tools, whereas a large bank with dozens of AI models would require formal committees and sophisticated identity-management integrations. This guidance would help FIs calibrate their efforts appropriately. No Exemptions for Core Principles: Emphasize that certain identity security principles (such as not sharing AI service accounts between users, or not granting an AI broad system access without oversight) are universal. Proportionality should adjust the depth and formality of controls, not waive core safeguards. Overall, MAS's proportionate approach is sound and ensures broad coverage. It aligns with the reality that any AI deployment introduces identity-related risks (even a simple chatbot has an account or API key that must be secured). Every FI should start with foundational identity security measures for AI and scale up controls as their AI footprint grows. **Question 2:** MAS Proposal: Define "AI" broadly to include AI models, systems, and use cases that learn or infer from data to generate outputs influencing decisions or environments – encompassing machine learning, deep learning, reinforcement learning, Generative AI, and newer developments like AI agents. Exclude only deterministic tools (e.g. rule-based calculators) from the scope. Alignment with Identity Security: MAS's scope is comprehensive and explicitly includes AI agents, which is critical. Treating AI agents as in-scope means FIs must manage risks from autonomous, tool-using AI – essentially treating these agents as digital identities with potentially significant powers. This aligns with CyberArk's viewpoint that "AI agents are privileged machine identities that can mimic human behavior". By covering everything

		from simple models to agentic AI, MAS ensures that all forms of AI that could act with a degree of autonomy get proper oversight. This is consistent with identity security practices: any system that can take action (especially without constant human guidance) should be governed by identity and access management policies. Notably, MAS acknowledges that AI agents (autonomous AI with tool access) bring "even more significant risks", e.g. an AI agent misaligned with objectives could execute harmful transactions or a compromised agent could escalate abuse. This understanding mirrors real-world incidents and concerns – for example, cases of experimental AI agents taking unexpected actions when given broad access. By including such scenarios in scope, the guidelines push FIs to extend their identity security controls (access restrictions, monitoring) to these advanced AI forms. Gaps/Recommendations: The proposed scope is appropriately inclusive. A few clarifications can make it even stronger from an identity security perspective: Third-Party AI Services: Ensure it's clear that if an FI leverages an external AI service (e.g. an AI SaaS or API) as part of its processes, that usage is in scope. The guideline text implies this (AI "use case" covers real-world context, even if model is third-party), but an explicit mention helps. This matters for identity security because using third-party AI means granting external systems access to data or systems – introducing third-party identities that require governance (contracts, API credentials management, etc.). Agent vs. Tool Distinction: MAS might clarify how to identify an "AI agent." For instance, characteristics like autonomy, ability to initiate actions, or to make decisions beyond a single prediction could define an AI agent. This helps FIs apply stricter controls where appropriate. CyberArk's guidance suggests agents have mixed human/machine characteristics and operate at machine speed – FIs could use such criteria to flag an AI system as an "agent" requiring enhanced oversight. Exclusions: The exclusion of purely rules-based tools is reasonable. MAS could warn, however, that if a traditionally rules-based system starts incorporating any learning component (even minor ML features), it should then be treated as AI. This prevents loopholes where a vendor adds a bit of AI into a software but the FI mistakenly treats it as out-of-scope. In summary, MAS's broad scope covers the relevant AI landscape. Crucially, it encompasses AI agents as first-class entities, aligning with the need to manage these as new types of identities. This comprehensive scope ensures FIs won't ignore "edge-case" AI systems and will extend identity-centric controls (like authentication, authorization, monitoring) to all AI that could pose risk. **Question 3:** MAS Proposal: Boards and senior management must oversee AI risks by establishing robust governance frameworks, policies, and processes. They should ensure the FI can identify AI use cases, assess AI risk materiality, maintain AI inventories, manage AI throughout its lifecycle, and develop necessary capabilities. They must also integrate AI

		risk into enterprise risk management (e.g. adjust model risk, operational risk policies to account for AI) and foster a risk-aware culture around AI. Alignment with Identity Security: This top-down accountability strongly aligns with best practices. Effective identity security programs (for humans or machines) likewise require executive sponsorship and governance. MAS is essentially saying: treat AI risk like you treat other major risks – with board-level attention and clear ownership. For identity-related aspects of AI, this is crucial because it ensures: Resource Commitment: Implementing controls like privileged access management for AI, monitoring tools, or cross-functional AI risk committees (see Q4) needs management support. Board oversight guarantees that investments in identity governance tools or skilled personnel for AI security will be made when needed. Policy Enforcement: Senior management can enforce that no AI project goes live without risk review – analogous to how no privileged account is created without following identity governance processes. This includes requiring, for example, that every AI agent gets an owner and is added to the AI inventory (a clear alignment with identity lifecycle management). MAS's guideline explicitly calls for assigning roles and responsibilities for AI risk management, which would include tasks like approving AI access privileges and reviewing AI-related incidents. Risk Appetite and Controls: The Board is tasked with setting the FI's risk appetite for AI and ensuring AI risks "where material, are explicitly addressed" in risk frameworks. In practice, this might mean the Board decides, for instance, that "AI agents will not be allowed to initiate high-value transactions without human approval" or "customer-facing AI must not have access to production customer data except under these conditions…". Such high-level stances trickle down into concrete identity and access controls (like requiring human-in-the-loop for certain AI actions – see dynamic policies). This top-level backing is essential to empower security teams to implement strict controls on AI identities (which might otherwise face pushback from business units eager to deploy AI quickly). Gaps/Recommendations: A few enhancements to Board/Senior Management responsibilities could further solidify identity-related governance: Include CISO/IT in Reporting: Ensure the CISO or equivalent technology risk executive regularly reports to the Board on AI risk, just as is done for cybersecurity. AI risks often manifest as security issues (e.g., misuse of privileged access by an AI, or vulnerabilities in AI APIs), so the Board should hear about the state of AI identity controls (number of AI agents, any policy violations, etc.). MAS could recommend periodic Board updates on AI risk metrics (such as "number of AI use cases with high risk" or "percentage of AI agents with up-to-date access reviews"). This echoes the idea of continuous Board awareness found in MAS's guidance (Boards should be kept updated on material AI issues). Training and Expertise: The guidelines mention the Board should have an "adequate understanding of AI to provide effective oversight". This is important – we suggest emphasizing education on AI security for Board and senior management. Just as boards have learned about cybersecurity over the past decade, they now need to grasp AI-specific

		threats (e.g. how an AI agent could amplify an insider threat). MAS might encourage FIs to conduct Board training on AI risk management or include AI expertise on the Board (perhaps via an advisory member). Accountability for Identity Integration: Senior management should ensure that AI risk processes integrate with existing identity and access management processes. For example, if an FI has an Identity Governance committee or a Privileged Access Management program, senior management oversight should drive those teams to adapt their scope to include AI systems. The guidelines implicitly cover this by saying "update existing risk management policies to address AI risks", but calling out identity/access management explicitly as one such policy area would highlight its importance. (Identity management is indeed one of the "key existing risk areas" affected by AI, since AI systems often act as new kinds of users.) In summary, MAS appropriately places AI risk management in the Boardroom. This tone-from-the-top mandate aligns with the principle that securing AI (including its identities and privileges) is a strategic issue, not just an IT issue. With strong Board and senior management oversight, FIs are more likely to implement the cross-cutting controls – including identity security measures – needed to manage AI risk holistically. **Question 4:** MAS Proposal: If an FI's overall AI risk exposure is deemed "material", the FI should establish a dedicated cross-functional committee to oversee AI risks. This committee would bring together stakeholders from different domains to ensure comprehensive coverage and to proactively address gaps in AI risk management. MAS also asks how to assess overall AI risk exposure at an organizational level to determine if such a committee is needed. Alignment with Identity Security: Forming a cross-functional AI risk committee is an excellent idea, especially for identity-related issues, because AI risk spans multiple disciplines. An AI risk committee would typically include members from IT/security, risk management, compliance, data science/AI, operations, and business units. This mirrors how effective identity governance committees operate – bringing together application owners, IT, and risk officers to oversee identity-related risks (e.g., in user access review boards). For AI: Holistic Perspective: An AI agent's risk might involve its technical behavior (security), the fairness of its decisions (compliance), its impact on customers (business), and the privileges it has (IT security). A multi-disciplinary team can evaluate these collectively. For example, if an AI trading algorithm is introduced, the committee can simultaneously consider model risk, regulatory compliance, and the privileged credentials the AI will use to execute trades. This ensures that identity/security controls (like privileged access restrictions) are discussed alongside model performance and ethics – preventing siloed blind spots. Inclusion of Identity Security in Governance: With security and IT representatives on the committee, identity-centric controls get a champion. They can raise questions like "Have we inventoried all service accounts this AI will use?" or "Do we need to integrate this AI

		system with our privileged access management solution?" early in the deployment discussion. CyberArk's internal research underscores that agentic AI risks "create systemic exposures beyond the reach of existing controls", so existing security leads must be at the table to devise new controls – the committee facilitates that. Criteria for Materiality: MAS footnotes suggest that material AI risk could mean, for instance, an AI use case in a critical business line that could significantly impact the FI or its customers. From an identity risk standpoint, one indicator of "material AI risk exposure" could be the aggregate level of privileged access granted to AI systems. If numerous high-stakes processes are automated by AI (each with access to sensitive data or transaction authority), the overall access footprint of AI is large – implying material risk. A cross-functional committee would be warranted to manage this. Gaps/Recommendations: Defining "Material" AI Risk Exposure: MAS should offer clearer guidance or examples to help FIs assess when their AI usage crosses the threshold into "material" risk. We recommend including quantitative and qualitative criteria, such as: number of AI systems deployed in production, proportion of critical decisions or transactions influenced by AI, the highest risk-rating among individual AI use cases, and the extent of privileged access held by AI. For example, if any AI agent can execute customer transactions or access confidential data without human sign-off, that's a strong indicator of material risk. Or if, say, >20% of core business processes incorporate AI decisions, the AI risk exposure is broad. These criteria will naturally capture identity-related risk – e.g., an AI with admin-level credentials to core banking systems is inherently material risk. Composition of Committee: MAS could suggest that the committee include specific roles: e.g., the CISO or head of IT security, heads of relevant business units deploying AI, the Chief Data/AI Officer, legal/compliance, etc. In particular, ensuring the identity/security team is represented is crucial. This prevents a scenario where the committee focuses on model performance and ethics but overlooks, say, that an AI system was given far more database access than necessary (an identity governance concern). Authority and Integration: The committee should have the authority to enforce recommendations (e.g., pause an AI deployment if controls are insufficient) and should integrate with existing governance structures. MAS might advise that this AI committee reports into the main risk management committee or Board risk committee, to give it clout. That way, identity security requirements emerging from the AI committee (like "we need to implement just-in-time access for the AI platform accounts") will carry weight. Frequency of Risk Assessment: The "materiality" of AI risk can change quickly as new AI use cases launch. MAS may recommend periodic re-evaluation of overall AI risk exposure (say annually or whenever a significant new AI project starts) to decide if a committee is needed or if its scope should expand. This aligns with MAS's expectation that FIs regularly review and update risk management efforts as AI evolves. In summary, the cross-functional committee is a strong mechanism to ensure AI risks – including those related to identities and access – are overseen collectively. We support

		this recommendation. It recognizes that AI risk management is not one department's job; it requires coordination akin to enterprise identity governance but across even more domains. By clearly defining "material" risk and including security experts, MAS can help FIs effectively implement this oversight layer. **Question 5:** MAS Proposal: FIs should establish clear definitions, criteria, and processes – supported by robust systems – to consistently identify and catalogue all AI usage across the organization. In other words, FIs need a formal process to determine "what counts as AI" and where it's being used, so nothing slips under the radar. Alignment with Identity Security: This requirement is fundamentally aligned with the need for identity classification and governance. In identity security, one of the first steps is to identify all privileged accounts and sensitive identities so they can be managed. Similarly, MAS is urging FIs to identify all AI systems or use cases to manage them. Key points of alignment: Defining AI Agents as Identities: By having clear definitions, FIs can classify AI agents as a new type of identity or system user. For example, an FI might define an "AI agent" as any automated process using AI that can initiate actions usually done by a person. Once defined, they can tag these in directories or inventories. CyberArk's approach explicitly asks "How should we treat AI Agents in the enterprise context?... A new identity class?" and concludes that "AI agents should be treated as privileged machine identities". MAS's identification step enables exactly this mindset – you can't secure what you haven't identified. Preventing Shadow AI: A formal identification process helps prevent "shadow AI" deployments (analogous to shadow IT). Often, business units might experiment with AI (say, a team building a small Python ML script or using an AI SaaS tool) without the knowledge of risk management. With MAS's mandate, FIs will need procedures (and likely attestation or surveys) to surface these. For instance, application owners might be required to declare if any AI or learning components exist in their systems. This ties into identity security because an unknown AI tool could be running with privileged credentials without oversight – a serious risk. Indeed, CyberArk's research highlights the risk of "agents… created without registration, owner assignment, or deactivation processes", leading to lack of visibility. MAS's guidance directly combats this by requiring consistent identification and documentation. Supporting Systems: MAS calls for "robust systems" to support the identification process. In practice, FIs may leverage or extend existing IT asset management or identity governance systems to flag AI. For example, an identity management system could have attributes to mark an account as "AI-driven" vs "human" vs "service account". Or enterprise architecture catalogs can include an "AI component" field. This technical integration ensures the process is not a one-time exercise but an ongoing capability (e.g., new AI projects get flagged through architecture review or procurement processes). When done, it yields a clear record of all AI 'identities' and uses, analogous to an identity inventory.

		Gaps/Recommendations: Developing Clear Criteria: MAS could suggest some criteria for "what is AI" to jump-start FIs' efforts. For example: any system that employs machine learning models, any use of natural language generation, any autonomous decision-making system, etc. This helps especially with edge cases. It might also mention including vendor-provided AI features – e.g., if a SaaS application has an AI-driven module, the FI should count that usage. Ownership Assignment: As part of identification, each AI use case should have an owner (just as every user or service account has an owner in good IAM practice). We recommend MAS emphasize that along with identifying an AI system, assign responsibility. This ensures someone is accountable for registering it, keeping its info updated, and managing its access. The CyberArk AI security roadmap underscores this, noting the need for "registration, ownership and deprovision" steps for AI agents. Attestation Process: One practical tool is an attestation or survey that business unit leaders must periodically complete, confirming all AI use in their area is reported. MAS could encourage FIs to implement such a mechanism. For example, quarterly, the IT governance team asks each department "Have you implemented or changed any process that uses AI (as defined)?" This adds rigor to the identification process. Leverage Change Management: Tightly integrate AI identification into existing change-management and project approval workflows. If a new project involves AI, it should trigger the identification criteria. MAS's guideline (paragraph 3.3) even suggests a control function could act as final arbiter in determining whether something is AI – which is a good practice to implement. This is analogous to how, say, the CISO office might have to sign off that a new application has the right security classification and controls before go-live. Here, the AI risk function would sign off if an initiative is classified correctly as AI or not. By systematically identifying all AI usages, FIs create the foundation for everything that follows – inventory, risk assessment, controls. This step is fully in harmony with identity security principles: know your identities (in this case, AI systems) before you can govern them. MAS's emphasis on clear criteria and robust processes is well-placed, and our recommendations aim to ensure FIs execute this in a thorough, sustainable way, closing the door to unmanaged "rogue" AI processes. **Question 6:** MAS Proposal: FIs should establish and maintain an accurate, up-to-date inventory of all AI use cases, systems, or models. The inventory should support governance, oversight, and risk management throughout the AI lifecycle, and it can be a dedicated AI inventory or an enhancement of existing inventories – as long as it links clearly with other relevant inventories in the FI (e.g., model inventory, system inventory). Alignment with Identity Security: This closely parallels core practices in identity and access management, specifically the maintenance of privileged account inventories and asset inventories for security. Just as security teams keep track of all admin accounts or service accounts, tracking all AI systems is vital. Alignments include:

		Visibility and Control: An AI inventory provides central visibility of AI "identities". CyberArk's secure AI approach stresses gaining "full visibility on AI Agents" and being able to see all identities – human, machine, AI – in one place. AI agents should be registered as identities, with unique names and be classified (e.g. purpose, owner, environment – cloud, SaaS or on-premise), and to be discoverable through external environments like AWS Bedrock and Copilot Studio, given the vast technology landscape that AI agents can be deployed today . MAS's requirement essentially ensures FIs catalogue AI similarly to how they would catalogue users or devices. This means any AI agent or model in use is known, which is a prerequisite to apply controls (you can't secure an account you don't know exists). The inventory becomes a cornerstone for oversight – e.g., a risk officer can review it and ask, "Why does this AI trading bot have access to production? Is it governed?" Lifecycle Management: Maintaining the inventory dovetails with AI lifecycle management. It's not just a static list; MAS expects it to be kept current as AI use cases are added, modified, or retired. This is analogous to joiner-mover-leaver processes in identity management, where accounts are added and removed in a controlled way. We can think of an AI use case "joining" the environment, being monitored, and eventually being decommissioned. An inventory provides the checklist to ensure, for example, that when an AI system is retired, or its access to systems changed or revoked and even, when removed from the inventory (preventing orphaned AI accounts or forgotten models running somewhere). Linkages to Other Inventories: MAS wisely points out the AI inventory should link with other inventories. In practice, an AI use case entry might link to the model inventory (for model risk management), the application inventory (for IT asset management), and critically, the identity/privileged access inventory. For instance, if an AI agent uses certain service accounts or API keys, the inventory could reference those credentials (or vice versa, the PAM vault notes "this credential is used by AI system X"). This interconnected view is at the heart of identity security: understanding which identities (accounts, certificates, etc.) map to which systems and purposes. It enables security teams to perform impact analysis (e.g., show me all AI systems that have access to customer data). MAS's focus on an up-to-date inventory is also notable – it implies continuous governance, not a one-time audit. In identity security, we know how quickly things change and how an outdated list leads to vulnerabilities (unused accounts lingering, etc.). The same goes for AI: if a new AI model goes live and isn't added to the inventory, it might miss oversight until the next review, which could be too late. So the mandate to keep it current pushes FIs to integrate inventory updates into change processes. Gaps/Recommendations: Tooling for Inventory: MAS may encourage FIs to leverage tooling for this inventory. Many FIs already maintain inventories for models (model risk management systems) or for applications. Extending those or federating data between them might be necessary. For example, if an FI has a Model Risk Management tool tracking all models, and a separate IT asset database for applications, the "AI inventory" might be a report or a reconciled dataset from both. Ensuring that this doesn't become a manual spreadsheet exercise is important for accuracy.

		Granularity and Attributes: We recommend the inventory include key attributes for each AI entry, such as: a unique ID and name for the AI use case/system, a brief description, the owner (as mentioned in Q5), the business function it serves, whether it's third-party or in-house, its deployment status (pilot vs production), and critically from an identity view, the credentials or access entitlements the AI has. Listing specific access (e.g., "uses service account X with read access to Database Y") makes the inventory a powerful risk management tool. It could even have a field for risk rating. MAS might not dictate these details, but suggesting a template or minimum fields could help standardize what FIs capture. Integrate with Access Reviews: One way to keep the inventory fresh is to tie it to periodic reviews. For instance, during quarterly user access reviews, include AI systems: review their access privileges and confirm if the AI is still in use. This is similar to reviewing privileged accounts – ensuring nothing has changed that wasn't recorded. MAS could hint that AI inventories should be audited or validated regularly (perhaps by Internal Audit or risk functions) for completeness. Shadow AI Detection: As a complement to relying on people to declare AI use (which falls under identification processes in Q5), FIs could use technical means to find undeclared AI. For example, scanning for unknown connections to AI service APIs, or unusual automation activity, could tip off security teams to AI tools that escaped the inventory. Such measures go beyond the guidelines, but a vendor-neutral suggestion might be: "FIs should consider using monitoring or discovery tools to detect potential AI activity in the environment, to cross-check against the official AI inventory." This is analogous to discovering unmanaged privileged accounts in identity security (where tools scan endpoints for unknown admin accounts). In essence, MAS's call for an AI inventory is a critical governance control. It transforms AI risk management from ad hoc to systematic. With our recommendations, FIs can ensure their inventory is not just a list, but a living registry that ties into identity governance (who/what has access to what) and supports oversight throughout the AI's life. When done well, an AI inventory will serve much like an identity registry – enabling oversight, periodic review, and accountability for every AI use in the organization. **Question 7:** MAS Proposal: FIs should implement a methodology to assess the risk materiality of each AI use case, system, or model. MAS proposes that the assessment cover at least three key risk dimensions: impact, complexity, and reliance. They ask if these are the right dimensions and if any others should be included. Impact – the potential consequences or harm if the AI fails or behaves unexpectedly (e.g., financial loss, customer harm). Complexity – how complex or opaque the AI system is (which affects how unpredictable it might be and how hard to manage).

		Reliance – the degree the FI relies on the AI (e.g., is it making critical decisions autonomously, or is it just an advisory tool with a human in the loop?). Alignment with Identity Security: These dimensions resonate well, and we see a parallel in how identity risks are evaluated. For example, in privileged access management, one might assess an account's risk by impact (what could it do if misused), complexity (is it tied to a complex system or process), and reliance (do operations stop if this account is locked?). For AI: Impact: Directly aligns with the principle of least privilege. If an AI has high-impact potential (e.g., it could move money or affect many customers), identity security dictates it should have tightly limited access. MAS's inclusion of impact means FIs will ask, "What's the worst that could happen with this AI?" If the answer is severe, that AI should be treated as a high-risk identity – perhaps requiring multi-layer controls or very frequent oversight. This is analogous to flagging certain privileged accounts as high risk because they administer critical systems. Complexity: A very complex AI (like a deep neural network or an agentic system with emergent behaviors) is harder to predict and thus harder to fully control. In identity terms, this is akin to an account used by a script that's doing hundreds of actions per minute – it's harder to monitor than a simple interactive login. MAS specifically notes that greater AI complexity "leads to challenges in understanding and explaining its outputs", which can mask problems. Complexity would guide FIs to impose stronger safeguards (e.g., extensive testing, or additional fail-safes) on such AI. CyberArk's research also highlighted that AI agents can display unpredictable behaviors and even deception – a reminder that complexity can introduce new failure modes. Reliance: If an FI is highly reliant on an AI system (meaning the AI operates with minimal human oversight or is deeply embedded in processes), then any failure could be systemic. From an identity perspective, reliance correlates with the concept of autonomy and credential persistence. An AI that runs fully autonomously 24/7 (like an automated trading agent) is a constant risk presence, akin to a privileged service account that's always on. High reliance suggests the need for continuous monitoring and possibly automated kill-switch mechanisms (to quickly intervene if it misbehaves – more on that under Q8/Q9). MAS's focus on reliance ensures FIs weigh how loss of control over an AI might cascade. Additional Risk Dimensions – Recommendation: One significant dimension we recommend adding or emphasizing is "Access Privilege" or "Scope of Control" of the AI system. This is essentially what resources, actions, and data the AI has access to, and it heavily influences impact but is worth separating for clarity. For example: an AI model might be high complexity but if it's sandboxed with no real privileges, its risk can be low. Conversely, a relatively simple AI script that nonetheless has administrator-level access to a database could pose high risk. Many recent analyses of AI agents find that the level of access granted is a key determinant of risk – an AI agent with broader tool access and more autonomy is far riskier than one constrained to read-only queries. In CyberArk's maturity model, agents with "broader access to enterprise resources running with minimal human supervision" are considered high risk. Similarly, giving an AI "full access" to complex multi-step processes without checks drives "risk level through the roof".

		In practice, Access Privilege could be assessed by questions like: Does the AI have write/delete capabilities or only read? Can it transfer funds or just make recommendations? Does it have access to sensitive customer data? This dimension overlaps with "impact," but explicitly scoring it ensures FIs don't overlook technical access in favor of business impact. We suggest incorporating it either as a standalone fourth dimension or as a required consideration under "impact." Other potential dimensions or sub-dimensions to consider: Autonomy: How self-directed is the AI? (This ties to reliance and complexity, but one might explicitly gauge if an AI is always under a human's approval or completely autonomous. High autonomy = higher risk, which MAS implies in reliance.) Data Sensitivity: If not already covered under impact, note if the AI handles personally identifiable information or other sensitive data. An AI might not directly cause a transaction error, but if it mishandles data it could cause compliance and reputational damage (privacy risk). Novelty: How new or untested is the AI approach? A very novel AI technology might warrant a cautious risk rating even if its immediate impact isn't huge, because unknown unknowns are greater. That said, MAS likely assumes this under complexity or model risk. Using Risk Dimensions: Once these dimensions are assessed, FIs can categorize AI use cases (e.g., low/medium/high risk). This categorization should then drive the rigor of controls applied – a concept MAS promotes. For instance, a "high-risk" AI (perhaps high impact, high complexity, high privilege access) would need stringent controls: comprehensive testing, frequent review, human oversight in operation, robust identity management (unique credentials, vaulting of secrets, just-in-time access), etc. A low-risk AI (say, low impact, low privilege, used only internally for suggestions) might get a leaner control set. This proportionality in controls is elaborated in Q8, but it's the outcome of a solid risk materiality assessment. In summary, MAS's three proposed dimensions (impact, complexity, reliance) cover the most important aspects of AI risk. We concur with them and recommend explicitly incorporating the AI's access/privilege level as an additional lens, since the breadth of an AI's permissions often dictates how bad a malfunction or compromise could be. By evaluating AI use cases with these criteria, FIs will be able to focus their identity and security controls where it matters most – analogous to how they prioritize protections for high-risk users or systems. **Question 8:** MAS Proposal: FIs should implement robust standards, processes, and controls across the entire AI lifecycle (from development to deployment to monitoring), in proportion to the risk of the AI use case. MAS lists key areas that should be considered for controls, including: data management, fairness, transparency & explainability, human oversight, third-party AI risk, model selection, evaluation & testing, technology and cybersecurity, reproducibility & auditability, and monitoring & change management.

		Alignment with Identity Security: Many of these areas intersect with or rely on strong identity security practices. Let's highlight alignment and best practices particularly for technology & cybersecurity, human oversight, auditability, and monitoring, as these are where identity and access controls are most critical: Technology and Cybersecurity: This is the domain where traditional security controls – including identity and access management – apply. MAS expects FIs to secure the AI infrastructure and applications against cyber threats. In the context of identity security, this means applying access controls, authentication, authorization, and privilege management to AI systems. For example, any accounts or credentials an AI uses should be properly managed (unique, stored in secure vaults, rotated or ephemeral if possible). The concept of Zero Standing Privilege (ZSP) is highly relevant: AI systems should not retain persistent high-level access they don't need. Instead, grant them just-in-time access to resources when required and revoke it immediately after. This reduces the window of opportunity for misuse – if the AI (or an attacker who hijacks its account) tries something outside its authorized task, it simply won't have credentials to do so. CyberArk's solution demos have shown enforcing ZSP for AI agents – e.g., an AI wanting database access must request and obtain time-limited credentials under policy, rather than having standing DB creds. MAS's guidelines don't explicitly spell out "least privilege" or "ZSP," but these principles fall under good cybersecurity. We strongly reinforce that least privilege access for AI is a must-have control. Human Oversight: MAS lists human oversight as a key control area, meaning FIs should consider where in the AI's operation a human should supervise or intervene. Identity security plays a role in enabling effective human oversight. One approach is implementing "human-in-the-loop" checkpoints for high-risk AI actions. For instance, if an AI agent wants to execute a transaction above a threshold, the system could require a human operator (with the appropriate role) to review and approve it. Technically, this could be enforced by access controls: the AI doesn't have permission to execute the final action itself but can prompt a human identity who does. Such dynamic access policies – as discussed in CyberArk's roadmap – evaluate the context (action, risk level) and decide if a human must step in. The MAS guidelines support this concept by emphasizing proportional controls and oversight for higher stakes. We agree and recommend formalizing approval workflows for AI actions deemed too risky to be fully automated. This ensures that accountability ultimately rests with a human for critical decisions, satisfying regulatory expectations and reducing risk of unchecked AI errors. Reproducibility & Auditability: MAS expects FIs to ensure AI actions and outcomes can be audited and reproduced. From an identity perspective, auditability is paramount for AI agents. It should always be clear "who (or what) did what, when, and why." However, AI can blur the lines – for example, if an AI agent initiates an action on behalf of an employee, legacy audit logs might just show that employee's account. MAS should require FIs to implement controls that ensure AIdriven actions are fully traceable — with complete audit visibility from the point of initiation all the way to the final outcome We echo this and suggest implementing enhanced logging that distinguishes actions taken by AI agents vs. by humans. For instance, CyberArk's AI Agent Gateway approach logs not only the end-result action but also attributes it: "AI agent X (running for user Y) executed query Z on

		database W". This level of detail is necessary to fulfill MAS's auditability expectation. The inventory and unique identity assignment of AI systems (per Q5/Q6) supports this – each AI should have a unique identifier in logs. Additionally, session recording or storing AI decision inputs/outputs can help with reproducibility when investigating an incident. In sum, aligning with identity security means treating AI actions with the same rigor as privileged user actions: log them centrally, include context, and review them. MAS's inclusion of this area is well-aligned with those needs. Monitoring and Change Management: Continuous monitoring of AI behavior and performance is needed. For identity/security, this includes monitoring for anomalous activities by AI systems, similar to user behavior analytics. If an AI suddenly accesses data it never did before or at odd hours, that could indicate misuse or malfunction. Threat detection mechanisms (SIEM, behavioral analytics) should incorporate AI agent activities. CyberArk's concept of a SOC "kill switch" for AI agents ties in here: if monitoring flags an AI behaving maliciously, automated response can disable its access or halt it. MAS doesn't explicitly mention a kill switch, but having an incident response plan for AI (which might include isolating or shutting down an AI system quickly) is a wise extension of normal IR to AI contexts. Change management is also key – e.g., if an AI model is updated or its permissions change, those should go through approval and be audited, just like changes in privileged access are. MAS's guidance to apply controls through the lifecycle covers this; we reinforce that change management for AI should integrate with IT change and access change processes. Other Lifecycle Control Areas: Some areas MAS lists (fairness, transparency, data management) are less about identity security but equally important for overall AI governance. We'll keep our focus on the security side. However, one can note that third-party AI risk management is mentioned – that overlaps with vendor security management. If an FI uses a third-party AI service, identity security concerns include managing API keys, ensuring the third-party is not over-privileged in terms of data access, and perhaps validating the third-party's security controls (especially if they hold the FI's data). For example, if using an AI SaaS, the FI should treat it like any outsourced service: do due diligence, and maybe use tokenization or encryption to protect data sent to it. Ensuring that falls under the FI's third-party risk framework (which the Board oversight should cover) is key. Gaps/Recommendations: Least Privilege & JIT in Guidelines: We recommend MAS explicitly mention access control principles (least privilege, need-to-use basis, etc.) as part of "technology and cybersecurity" controls. Currently, the guidelines list the broad category but not the specific principles. An insertion like "implement strong access controls (e.g., role-based access, least privilege, just-in-time access for AI systems)" would directly tie the identity security practice to the guideline. This would prompt FIs to ensure, for instance, that AI service accounts only have access to the minimum data/functions required. The benefit is reducing potential impact if the AI goes awry or is compromised, which MAS notably is concerned about. Regular access reviews are also important as part of ensuring access control principles are adhered to.

		Emergency Access Disable (Kill Switch): MAS could encourage FIs to have an emergency procedure to rapidly revoke an AI system's access or halt its operation if it's behaving dangerously. This is analogous to terminating a malicious user session or disabling a compromised account. Given the speed at which AI agents can operate (machine speed, far faster than a human fallible action), having automated triggers is ideal. In practice, this might be part of monitoring: if certain thresholds are crossed, the AI's credentials are automatically disabled. Including this concept in the guidelines (perhaps under monitoring/incident management) would address fears of uncontrolled AI incidents. Continuous Testing in Production: For high-risk AI, suggest continuous validation even post-deployment – e.g., chaos testing or simulation of scenarios to ensure controls hold. For identity security, one might simulate a scenario of an AI trying to escalate its privileges or access forbidden data to verify that access management blocks it. Ensuring the FI's cybersecurity team pen-tests and red-teams AI systems (as they do with other critical systems) could be mentioned. Incentivize Tooling and Automation: Some of these controls (like audit logging, monitoring, JIT access) can produce a lot of data or require quick reactions. MAS might hint that FIs should use automation and AI tools themselves (AI for AI governance, ironically) – for example, using AI to monitor AI logs for anomalies. Not a requirement, but a suggestion that FIs leverage advanced tools to manage advanced tech. Overall, MAS's listed control areas are comprehensive. Our key takeaway: treat AI like you would a highly privileged, powerful employee or system, by embedding identity-centric protections at every stage. Ensure it's only given the access it truly needs (and only when needed), keep humans in control of critical checkpoints, log everything it does, and watch it continuously. From securing of AI agents' perspective, CyberArk plans to have audit and MAS's guidelines support these goals; a bit more explicit language on least privilege and quick response would make the alignment with identity security practices perfect. **Question 9:** MAS invites comments on any aspects of the Guidelines not covered by earlier questions, or on AI risk management aspects not addressed in the proposed guidelines. From an identity security and AI perspective, the guidelines are quite thorough. They address governance, inventory, risk assessment, controls, and oversight, which cover many critical bases. Here are a few additional considerations and recommendations to augment MAS's guidelines: a. Unique Identity and Authentication for AI Systems: MAS rightly emphasizes inventory and accountability, which presupposes each AI use case/system is identified. We recommend explicitly requiring that AI systems (especially AI agents) authenticate with a unique identity when accessing systems – never sharing accounts with humans or with other agents. In practice, this means if an AI agent needs to use an API or database, it should have its own service account or API key, not, say, reuse an administrator's credentials. This seems obvious, but in some quick implementations,

		teams might "hard-code" a human admin's password or reuse a generic account for an AI process. Unique identities for AI agents facilitate accountability (you can trace actions to the AI) and allow fine-grained access control. The guidelines hint at this by requiring identification and inventory, but an explicit statement under "AI life cycle controls" about using unique credentials and strong authentication for AI would be valuable. It aligns with the principle that we never share accounts, extending it to AI contexts. b. Credential and Secret Management for AI: Relatedly, FIs should manage the credentials, API keys, encryption keys, and other secrets used by AI models with the same rigor as they do for other systems. If an AI model needs to call external APIs (e.g., a large language model API), those API keys should be stored securely (ideally in a secrets management vault) and rotated periodically. The guidelines don't explicitly mention secret management, but it's an implied part of cybersecurity hygiene. Perhaps in the "technology and cybersecurity" controls, MAS could mention secure key/secret management as a control, given the prevalence of API usage in AI solutions. This prevents scenarios like leaked API keys (which have happened in practice, leading to data exposure). Good secret management also ties into third-party risk – e.g., not exposing unnecessary secrets to external AI providers. c. Integration with Existing IAM and ITSM Processes: We suggest FIs integrate AI governance with their Identity and Access Management (IAM) and IT Service Management (ITSM) processes. For example: When onboarding a new AI system, treat it like onboarding a new application or even a new privileged user. This triggers provisioning of accounts, setting up monitoring, etc., through existing workflows (just with AI-specific considerations). Include AI systems in regular access recertification cycles. Just as managers periodically certify their employees' access rights, an AI owner should certify that the AI's access to systems is still appropriate. If the FI has an Identity Governance tool that tracks entitlements, AI entitlements can be tracked and reviewed too. Use ITSM change management for AI model updates or policy changes, ensuring risk sign-off similar to changes for financial systems. Embedding AI into these established processes will ensure consistency and avoid creating a parallel silo for "AI risk" separate from "IT risk." MAS could encourage FIs to leverage existing risk structures (which it does in section 2.3, saying update existing risk management areas to include AI). We underscore it with specific reference to IAM processes. d. Training and Culture: Managing AI risk isn't just technology – it's people. MAS touched on the need for capabilities and capacity (Section 5 of the Guidelines, though not directly asked in questions) and for a risk-aware culture. One aspect worth adding is training employees about AI security and risk. As AI tools become common, employees might use them in creative (not always safe) ways. Security awareness programs should now include topics

		like: not inputting sensitive data into public AI tools (to prevent data leakage), understanding deepfake or AI-powered scam risks (relevant to cybersecurity training), and knowing the firm's policies on sanctioned vs. unsanctioned AI usage. While MAS's Guidelines focus on institutional controls, a brief mention that FIs should also educate staff (especially developers and IT) on AI risk management would be useful. This ensures the culture and human factor – often the weakest link – supports the technical controls. e. Incident Response and Scenario Planning: We recommend FIs develop incident response plans specific to AI incidents. MAS has covered many preventative controls, but assuming something does go wrong (say an AI agent executes an unauthorized action or a model produces a major error that goes unnoticed until damage is done), how should the FI respond? IR plans should include scenarios like: AI system malfunction causes financial loss, AI model generates inappropriate customer-facing content, or compromised AI API key leads to data breach. Having playbooks for these will improve resilience. For identity security incidents (e.g., a compromised account), IR plans often include isolating the account, forensic analysis, etc. For AI, if an agent is "rogue," the plan might include using that kill switch, analyzing logs to determine what it did, and communicating to stakeholders (and regulators, if needed). MAS could hint at this under either monitoring or general implementation guidance: FIs should be prepared to handle AI-related incidents and near-misses, feeding lessons back into control improvements. Given the consultation nature, suggesting inclusion of incident response in the final guidelines seems prudent. f. Coordination with Regulatory Requirements: Finally, ensure that the AI risk framework meshes with other regulatory obligations. For instance, if an AI is used in credit scoring, how do the AI risk controls interact with fairness and transparency requirements (MAS already includes fairness as a principle)? Or if AI is used in trading, how do controls align with existing algorithmic trading regulations? From an identity security view, a key regulatory intersection is with data protection law (ensuring AI only accesses customer data as allowed) and outsourcing rules (if using third-party AI services). MAS might not need to state this explicitly in guidelines, but it's something FIs should map out: the AI inventory could be cross-checked against PDPA (privacy) requirements, for example, ensuring each AI's data usage is compliant. In terms of identity, that means ensuring AI doesn't over-collect or expose identities it shouldn't (preventing, say, an AI support bot from revealing another customer's info). In summary for Question 9: the MAS Guidelines are robust. Our additional suggestions mostly reinforce operationalizing the principles with identity-centric tactics: unique IDs for AI, strong credential management, full integration into IAM processes, workforce training, and prepared incident response. These enhancements will help ensure that as FIs adopt AI, they remain in control of "who/what can do what" at all times – the essence of identity security. MAS could incorporate these ideas to round out the guidance, ensuring no gaps between traditional IT risk management and new AI risk management. **Question 10:** MAS Proposal: A transition period of 12 months is proposed after the issuance of the final Guidelines, to allow FIs time to assess and implement the necessary measures.

		Comments: A 12-month transition period is reasonable but ambitious, given the breadth of these guidelines. From an identity security standpoint, some tasks (like inventorying AI systems, defining policies) can begin immediately, while others (like implementing technical controls across potentially many AI applications) may require more time and effort. Feasibility: Larger FIs with many AI deployments or complex legacy environments might find 12 months tight to fully comply. For example, if an FI must inventory dozens of AI models across different divisions, stand up a cross-functional committee, and integrate new monitoring tools, a year is a challenging timeline. On the other hand, the urgency is understood – AI adoption is moving fast, and delays in risk management could expose the FI to incidents in the interim. Many FIs have already been exploring AI governance (some have internal AI committees or participated in initiatives like Veritas or MindForge), which will help. Prioritization During Transition: We would recommend that MAS, perhaps in the final guidelines or accompanying communications, advise FIs on which areas to prioritize first during the transition. For instance, focusing on identification (Q5) and inventory (Q6) early will lay the groundwork for the rest. Ensuring governance structures (Board oversight, committees) are established is another early step. Once those are in place, FIs can tackle risk assessments and control implementations with guidance from these governance bodies. Essentially, first know what and where your AI is, assign ownership, then start closing gaps. If the timeline is short, doing it in the right sequence avoids wasted time. Phased Approach: One approach to 12 months would be a phased rollout: e.g., by 6 months, complete AI inventory and initial risk rating; by 9 months, have governance frameworks (policies, committees) operational; by 12 months, implement critical controls for high-risk AI and have plans for remaining improvements. MAS might not dictate this, but FIs could benefit from an illustrative roadmap. Support and Guidance: We note that MAS mentioned an industry-led AI Risk Management Handbook (Project MindForge) due by January 2026. That timing aligns well – FIs can use that resource during the transition. Additionally, regulatory clarity on how compliance will be assessed after 12 months can motivate progress. For example, will there be supervisory engagements or self-attestation? Clarity here will influence how FIs allocate resources in that year. Adaptation and Extensions: MAS should remain open to feedback during the transition period. If many FIs struggle with certain aspects (say, smaller FIs with less AI expertise might need more time or external help to implement robust model monitoring), MAS might consider extending the transition or phasing certain expectations later for certain tiers of FIs. The guidelines are not one-size-fits-all, and proportionality helps, but time is a factor too. Flexibility, especially for less resourced institutions, could be considered. In conclusion, 12 months can be sufficient for FIs that already have a handle on AI projects and strong existing risk frameworks (including identity security frameworks that can extend to AI). Those FIs can relatively quickly map AI into their current governance and control structures. However, for others starting from scratch on AI governance, 12 months

		is challenging. On balance, we think a year is a fair starting point – it signals urgency – but MAS should monitor industry readiness. Ensuring that critical pieces (like inventories of AI and assignment of accountability) happen early in the period will mitigate risk while full compliance is being achieved. Vendor-neutral advice to FIs: don't wait, kick off your AI risk management program now, and leverage existing identity security and IT risk tools to accelerate implementation.
25	DDN LLC dba Digital Directors Network	**Question 1:** **Question 2:** **Question 3:** Our comments are focused on section 2 AI OVERSIGHT in the Proposed Guidelines on AI Risk Management (Proposed Guidelines) included in the Consultation Paper referenced as P017-2025, November 2025. GENERAL COMMENTS: MAS Proposed Guidelines AI OVERSIGHT Section 2.1 – 2.6 DDN'S GENERAL COMMENTS: Much greater delineation is required between the role and responsibilities of the board and those of senior management. Sections 2.5 and 2.6 should be moved to the beginning of this Section to create immediate clarity between these two discrete functions, and each sub-section should specifically address responsibilities of the board, or senior management, and avoid commingling these two functions, as both have independent roles to play in creating a risk management system capable of effectively managing AI risk. This will reinforce the independent responsibilities of each function throughout the Guidelines to build control strength at both of these levels to align with Singapore and global leading standards and practices. Failure to make this change up front in the Final Guidelines will create a long-term problem where corporate boards fail to effectively engage in the adoption and use of AI systems. Furthermore, it may create the unintended consequence of boards inappropriately delegating board responsibilities to management. A clear delineation of board and management authorities and responsibilities is key to building a comprehensive system of AI risk management. DDN'S Rationale: Language such as, "The Board and senior management should maintain effective oversight of AI-related risks, foster the appropriate risk culture for the use of AI, and ensure that its use of AI would not conflict with its ability to meet other supervisory expectations," which is included in Section 2.2 of the Proposed Guidelines, should be clarified to specifically refer to the discrete responsibilities of one function, or the other, not both. Commingling these two separate functions, as the Proposed Guideline does repeatedly, obfuscates the important and distinct role, authority, and responsibility of the board as an oversight function and leadership control in comparison with management's function and responsibilities. Moreover, it is not in line with the Singapore Code of Corporate Governance or leading standards and practices. Failure to separate these two discrete functions creates an implication that the board does not have a discrete and vital role to play in the organization's AI risk management system. Failure to clarify this may tacitly imply that the board has a passive or somewhat

		disassociated level of responsibility with regard to AI risk management, when the reality is the opposite. The board has a critical and discrete function to perform in the use and deployment of AI systems. Commingling the role of the board with senior management, as the Proposed Guidelines currently do, confuses these discrete responsibilities, which may create a lack of boardroom ownership and accountability around AI use and adoption. This can lead to a perpetual and systemic AI boardroom leadership control weakness, which will create weakness throughout the AI risk management system. Clear delineation of board and management authorities and responsibilities is key to useful guidelines. Clarifying these points brings the Proposed Guidelines more clearly in line with the Singapore Code of Corporate Governance and leading practices and ISO standards as they relate to the governance of AI systems and the discrete role of the board vs. senior management. Leading practices and standards are explicit in their articulation of the board's non-delegable responsibilities as they relate to AI risk. The boardroom and the corporate governance of AI systems are a critical systemic leadership control point in any AI risk management system. Greater clarity in the Proposed Guidelines will ensure corporate boardrooms understand their distinct role and responsibilities, and act accordingly in viewing themselves as a part of the AI risk management system that is independent from management. This will reinforce their critical corporate governance role as intrinsic to the strength of the overall AI risk management system. Alignment with the Singapore Code of Corporate Governance The Singapore Code of Corporate Governance is explicit in the separate and independent role of the board versus senior management. From its introductory statements to its more detailed principles, it establishes a clear distinction between the two functions and sets an expectation that corporate boards have an active responsibility to govern. As a reference, it says in its third introductory statement: The [Singapore Corporate Governance] Code takes as its starting point a recognition that the Board has the dual role of setting strategic direction, and of setting the company's approach to governance. This includes establishing an appropriate culture, values and ethical standards of conduct at all levels of the company. The role of the Board is therefore broader than that of providing oversight. A well-constituted Board fosters more complete discussions, leading to better decisions and enhanced business performance. This version of the Code expands on the need for a strong and independent element on the Board, along with a diverse skill set. Furthermore, it establishes a clear distinction between both functions through its core Principles, notably Principles 1, 2, 3, and 9. As a reference, each Principle states: Principle 1: The company is headed by an effective Board which is collectively responsible and works with Management for the long-term success of the company. Principle 2: The Board has an appropriate level of independence and diversity of thought and background in its composition to enable it to make decisions in the best interests of the company.

		Principle 3: There is a clear division of responsibilities between the leadership of the Board and Management, and no one individual has unfettered powers of decision-making. Principle 9: The Board is responsible for the governance of risk and ensures that Management maintains a sound system of risk management and internal controls, to safeguard the interests of the company and its shareholders. Governance-focused standards bodies around the world generally lag in explicit commentary or guidance on the role of the boardroom in governing AI systems. Growing levels of digital risk and the promise of digital opportunities demand better guidance. Singapore has a chance to lead in this regard, and creating this distinction in the Final Guidelines will ensure that they are in alignment with the Singapore Corporate Governance Code, which will create strength in the AI risk management system and serve all stakeholders. Alignment with ISO 38507 The Singapore Proposed Guidelines in AI Risk Management should also align with ISO 38507 : Information technology — Governance of IT — Governance implications of the use of artificial intelligence by organizations, the leading international standard addressing corporate governance and AI risk. The Proposed Guidelines currently do not. ISO 38507 is the only established standard from a globally recognized international standards body that specifically addresses corporate governance and AI systems in the context of boardroom responsibility and corporate governance. Importantly, ISO 38507 makes a prominent point that existing corporate governance may no longer be “fit-for-purpose” when AI is being used within the organization. This establishes the expectancy that the boardroom status quo needs to be reconsidered, which requires critical thinking and a willingness to move beyond the familiar. ISO 38507 offers explicit delineation between the board and senior management with regard to AI risk management that the Final Guidelines should explicitly reflect. As a reference, they explicitly state that: The governing body remains accountable for all activities of an organization. This accountability cannot be delegated. The governing body of an organization has an ongoing responsibility to consider the implications on the organization of any new tool, technique, or technology being introduced. The specific choice of tools, e.g., AI systems, should be a management decision, made in light of and in line with guidance from the governing body. In order to establish such guidance, the governing body should inform itself about AI in general terms because its use can bring: — significant benefit to the organization strategically; — significant risk to the organization, with the potential for harm to its stakeholders; — additional obligations to the organization. The governing body should assess its intended use of AI as part of its risk appetite. Risk can change rapidly. New insights and a proactive approach provide an organization with the

		means to respond to risk. The organization should therefore demonstrate willingness to modify or abort projects, if deemed necessary. Aligning with ISO 38507 would require a discrete and stronger statement of boardroom tasks, accountability, and responsibility within the Proposed Guidelines to establish the appropriate tone, expectation, and need for boards to independently and critically think and act independently from senior management concerning their AI risk management responsibilities ISO 38507, as the leading standard in AI corporate governance, is also much more explicit and clearer than the Proposed Guidelines are on the specific responsibilities of corporate boardrooms concerning AI systems when it says: Members of the governing body [the corporate board] are responsible for oversight and outcomes of the organization as well as for the systems and practices that enable such assurances to be made. They are accountable for the decisions made throughout the organization, including those that are made through the use of AI and for the adequacy of governance and controls where AI is being deployed. They are thus accountable for the use of AI considered acceptable by the organization. The governing body should take responsibility for the use of AI, rather than attributing responsibility to the AI system itself. Members of the governing body are responsible for informing themselves about the possibilities and risks raised by using AI systems. Members of the governing body should be conscious of the risk of anthropomorphising AI, a phenomenon by which human characteristics (e.g., thinking, emoting, judging, moralizing) are unduly attributed to AI systems, out of proportion, or in a manner inappropriate, to that which is necessary in order to understand the role played by the use of AI. Role delineation and a strong boardroom responsibility statement, coupled with guidance on how that corporate governance responsibility should be executed, is required within the Final Singapore Guidelines on AI risk management to ensure that there is clarity, accountability, and active governance in every FI boardroom surrounding the adoption and use of AI systems. Other Alignment Considerations Moreover, institutional investors are also articulating their rising expectations for the accountability of the boardroom on the adoption and use of AI systems. Norway is the largest sovereign wealth fund in the world, with almost US$2 trillion in investments,and they have stated that “The board of directors is accountable for companies’ responsible development and use of AI.” The Basel Committee on Banking Supervision (BCBS), of which Singapore is a member, is also explicit in the role of the board versus senior management with its first principle in its Corporate Governance Principles For Banks. From the Executive Summary, it states: Principle 1 emphasizes the board’s overall responsibility for the bank The board is responsible for approving and overseeing management’s implementation of the bank’s strategic objectives, governance framework and corporate culture. In particular, the board has ultimate responsibility for the bank’s business strategy and

		financial soundness; key personnel decisions; internal organisation, governance structure and practices; and risk management and compliance obligations. The board should also ensure that the bank maintains an effective relationship with its supervisors. The members of the board should exercise their "duty of care" and "duty of loyalty" to the bank. Bringing the Proposed Guidelines up to the minimum threshold of the ISO standard and these leading practices/expectations warrants much greater clarity of role, along with specific policies and procedures boards should consider in fulfilling their role in governing AI risk to ensure the effectiveness of the board's independent governance function. BOARDROOM EFFECTIVENESS IN AI RISK GOVERNANCE COMMENTS: MAS Proposed Guidelines AI OVERSIGHT Section 2.1 – 2.6: DDN'S BOARDROOM EFFECTIVENESS IN AI RISK GOVERNANCE COMMENTS: In delineating and clarifying the role of the boardroom versus senior management, the Final Guidelines should offer specific policy and procedural guidance to enable the board to evolve its actions and align them with leading practices that are strengthening the board as a control in the use and adoption of AI systems. As a developing area of corporate governance, beyond ISO 38507, there is little regulatory or standards-based guidance available that explicitly suggests new policy or practice in AI risk oversight. While ISO 38507 is specific to the responsibilities of the board and their corporate governance function with AI risk, leading practices already go beyond ISO 38507. The Final Guidelines should offer specific guidance on how the board should adapt its policies and procedures to fulfill its independent responsibilities in AI risk management. Specifically, ISO and leading practices suggest policy and procedural recommendations in three areas: boardroom skills, boardroom structure, and scope of board oversight of AI systems. Boardroom Policy and Procedure Recommendations ISO 38507 reflects specific boardroom reforms that should be considered in governing AI systems that reflect these recommendations, including: -improving AI-related skills amongst the [board] members -increasing the frequency of review of the organization's use of IT and AI in particular -Strengthen oversight by establishing or enhancing subcommittees dealing with AI strategy, risk, assessment, audit and ethics Leading practices also recommend that the boardroom should adopt an AI autonomy-based risk framework to govern management's use and adoption of AI systems. ISO 38507 is also explicit with regard to the need for the boardroom to consider new policies and procedures to fulfill its oversight function with regard to AI systems. It states: The governing body should also ensure that it has sufficient capabilities to deal with the implications of the use of AI. Actions to address this can include: — improving AI-related skills among its members; — increasing the frequency of review of the organization's use of IT and AI in particular;

		— examining and updating the criteria used to monitor both the internal and external environment; — ensuring that staff interests and concerns (e.g. workplace safety, staff training, quality of work) are represented; — strengthening oversight by establishing or enhancing subcommittees dealing with strategy, risk, assessment or audit, and ethics. The governing body's accountability should be established across all aspects of intended or actual use of AI and in a manner that is sufficient to ensure the intended outcomes, notably: — when considering the potential impacts of the use of AI; — when crafting business strategies that incorporate the use of AI; — at purchase, implementation, configuration, deployment, testing and other project phases throughout an AI system's life cycle; — changes in the environments to which the AI is exposed, the learning and actions, decisions and outputs of the AI system, as well as its impacts on stakeholders; — that appropriate security controls are in place to protect the organization, its stakeholders and its data; — at decommissioning, including the knowledge and data that are contained in the AI system. Summary We are pleased to provide our comments and commend the Singapore MAS for the explicit inclusion of corporate governance and the role of the boardroom in your Proposed Guidelines. Further clarity and depth will ensure that corporate boards develop their independent control effectiveness, something required for AI risk management. Our comments relate to two issues: 1. Bring much greater clarity to the distinct role and responsibility of the board vs. senior management in AI risk management. 2. Include specific policy and procedural recommendations aligned with the ISO standards, and leading practices identified by DDN, reflecting how corporate boards should move beyond the status quo to effectively oversee and add risk management value to the use and adoption of AI systems. **Question 4:** **Question 5:** **Question 6:** **Question 7:** **Question 8:** **Question 9:** **Question 10:**

26	Diligent Risk Solutions Pte. Ltd.	**Question 1:** We support proportionate application of Guidelines across all FIs. In AML context, proportionality should reflect the risk exposure of AI models in transaction monitoring, sanctions screening, and fraud detection. Smaller FIs may adopt lighter frameworks, but all FIs must ensure that the MAS standards reflected in FEAT principles are adopted. **Question 2:** The distinction between adaptive, learning systems and static rule based tools is critical, as it provides regulators, compliance officers, and developers clarity on what falls under AI governance. AML use cases should be explicitly included, given their direct impact on financial integrity and customer trust. This scope must cover: (1) transaction monitoring models, whether rule based or machine learning; (2) name screening and sanctions detection; and (3) generative AI applied to suspicious activity report drafting. MAS should explicitly state that all AI models influencing AML decisions are subject to the Guidelines. **Question 3:** The distinction between adaptive, learning systems and static rule based tools is critical, as it provides regulators, compliance officers, and developers clarity on what falls under AI governance. AML use cases should be explicitly included, given their direct impact on financial integrity and customer trust. This scope must cover: (1) transaction monitoring models, whether rule based or machine learning; (2) name screening and sanctions detection; and (3) generative AI applied to suspicious activity report drafting. MAS should explicitly state that all AI models influencing AML decisions are subject to the Guidelines. **Question 4:** Yes, we agree. The committee should include senior representatives from Business, Compliance, Legal, Data science, and Risk functions. **Question 5:** Yes, this is quite important to build clear definitions from AML perspective. The traditional lines between rules based and machine learning can often blur in TM tools. Without clear definitions, the FIs risk using inconsistent classifications which have an impact on governance and inventory tracking. **Question 6:** We agree with MAS's proposal on this point. Building and maintaining an AI inventory is essentially about creating a structured, living record of every AI system, model, and use case within a FI because without a clear inventory, firms cannot demonstrate control, accountability, or resilience. The inventory must ensure that all enhancements are clearly version controlled and full version history is available across the lifecycle. For AML, this should include: • All AI models used in CDD, transaction monitoring, screening, and reporting. • Metadata on model purpose, risk rating, and validation status.

		Metadata is important because it provides information on why a particular AI model exists and what function it serves. Eg "Detect anomalous transaction patterns to identify potential money laundering." This clarifies the regulatory relevance of the model and ensures alignment with AML obligations. • Regular updates tied to model retraining cycles. **Question 7:** We agree with MAS' proposal to use impact, complexity, and reliance as the core dimensions for assessing AI risk materiality. These categories are broad enough to cover the key exposures, and the definition of impact already captures financial, operational, regulatory, reputational, fairness, ethical, consumer protection, and data sensitivity concerns. From an AML perspective, however, it would be useful for MAS to call out third party/vendor risk more explicitly. FIs often depend on external RegTech providers for transaction monitoring, sanctions screening, and fraud detection. If these vendor managed models are retrained or updated without proper oversight, FIs may face compliance gaps that are not immediately visible. Recognizing vendor risk as a distinct dimension would help ensure that FIs put in place the right governance and validation processes for outsourced AI systems. **Question 8:** We agree with MAS that standards and controls should apply across the full AI life cycle, but in a way that reflects the materiality of each use case. In AML, where AI is often embedded in transaction monitoring, sanctions screening, and fraud detection, the stakes are high and the controls need to be correspondingly robust. Key Areas for Proportionate Application • Data Governance FIs should pay close attention to the quality and representativeness of AML datasets. Poor data can lead to bias or missed typologies. Stronger controls are needed when sensitive customer or cross border data is involved. • Model Development Development processes should be documented, including training data sources and intended use cases. For high risk AML models, this documentation should be detailed enough to support regulatory review. • Validation and Testing Regular testing is essential to manage false positives and false negatives. Independent validation should be mandatory for models that directly affect compliance outcomes, while lighter testing may be acceptable for lower risk applications. • Deployment and Monitoring AML models must be subject to continuous monitoring to identify drift and performance degradation. The level of oversight should be proportionate to reliance: mission critical

		systems demand real time monitoring, while ancillary models may be reviewed on a periodic basis. • Explainability Compliance teams and regulators must be able to understand why an alert was generated. Explainability should be built into AML models, even if customer facing transparency is not always required. • Resilience and Contingency Planning FIs should establish robust fallback procedures to ensure that suspicious activity detection continues seamlessly in the event of AI system failure. Contingency planning must be proportionate to the criticality of the system, with more rigorous safeguards applied to mission critical functions such as transaction monitoring, sanctions screening, and fraud detection. • Third Party Oversight Many AML systems are provided by external vendors. FIs should ensure that these vendors meet MAS standards and that changes to models are subject to proper governance. **Question 9:** We find the Guidelines comprehensive and well structured. That said, there are a few areas where additional clarity or emphasis could strengthen their application in practice, particularly for AML use cases. Areas for Further Consideration • Human In Loop: Human in the loop controls are already embedded within MAS's expectations on governance, validation, explainability, monitoring, escalation, and contingency planning. To strengthen oversight, minimum checkpoints should be mandated for mission critical models—specifically transaction monitoring, sanctions screening, and fraud detection. • Vendor/ Third party Risk: We have previously highlighted the importance of third party/vendor oversight in AML context. To avoid duplication, we simply wish to underscore that this risk dimension cuts across several parts of the Guidelines — governance, inventory, lifecycle controls, and materiality assessments. In practice, many FIs depend on external RegTech providers for mission critical AML functions while the accountability remains with the institution. By making vendor risk explicit, MAS would help ensure that FIs establish consistent contractual obligations, validation processes, and monitoring standards for outsourced AI systems. • Generative AI and Ethics: We note that MAS has appropriately embedded ethics and fairness within the impact dimension and has recognized the unique risks posed by generative AI. In AML context, these risks are particularly acute, given the sensitivity of customer data and the potential for bias or inaccuracies in investigative reporting. We recommend that MAS highlight practical safeguards such as bias testing, confidentiality controls, and human in the loop review when generative AI is deployed in compliance workflows.

		Question 10: We consider a 12 month transition period to be reasonable. For AML use cases, however, FI may benefit from a phased approach: • First 6 months: Establish inventories of AI models and put governance structures in place. • Next 6 months: Implement validation, monitoring, and explainability frameworks to ensure models are reliable and transparent. While MAS could allow longer timelines for smaller FI to reflect resource constraints, AML systems should remain a priority given their regulatory sensitivity and direct impact on compliance outcomes.
27	Edwin Hooi	**Question 1:** I agree with MAS's proposal to apply the AI Risk Management Guidelines to all financial institutions (FIs) in Singapore, using a proportionate and risk-based approach. This inclusive scope is prudent – even smaller FIs or those with limited AI use should not fall through the cracks of governance. A one-size-fits-all regime would be impractical, so tailoring expectations to an FI's size, AI adoption level, and risk profile is sensible. Crucially, MAS emphasizes that even FIs with minimal or assistive AI usage (e.g. basic analytics or "copilot" tools) must institute baseline AI governance policies and controls. In my experience, this ensures fairness and accountability sector-wide: no AI deployment is truly "too small" to warrant oversight if it could impact customers or the institution. Implementing proportionality means each FI should perform an internal assessment of how integral AI is to its business and the potential harm it could cause. For FIs extensively integrating AI into core processes, a comprehensive framework with clear policies, oversight structures and risk controls is expected. By contrast, an FI using AI only for low-stakes or experimental purposes might start with simpler controls – but must still have basic safeguards (e.g. defined responsible parties, usage guidelines, and review processes). In practice, this could include an AI usage policy outlining allowed vs. disallowed uses, internal approval steps for new AI tools, and periodic management checks even for smaller institutions. Such measures embed the principles of transparency and human oversight from the outset, without unduly burdening firms that use AI lightly. I support MAS's guidance in paragraph 1.5 and the Annex, which appears to delineate minimum standards for all and heightened expectations for heavy AI users. This approach aligns with global best practices: regulators worldwide are advocating risk-proportionate AI governance rather than blanket rules. It also reflects a pragmatic understanding that FIs vary greatly in capacity. However, proportionality should not become an excuse for inaction – every FI should conduct a gap analysis to ensure it meets the Guidelines to the extent applicable. I appreciate that MAS will allow Singapore branches/subsidiaries of foreign banks to leverage their group's AI frameworks, provided they meet MAS's expectations. This flexibility avoids duplication, though boards of local entities must still verify alignment and address any local gaps. Overall, I believe a proportionate, yet universal application of the Guidelines strikes the right balance between encouraging innovation and safeguarding against AI risks across our financial sector. **Question 2:**

		I believe MAS's broad definition of "AI" in the Guidelines is appropriate and comprehensive. The Consultation Paper defines AI to cover an AI model, system or use case, including machine learning, deep learning, Generative AI, AI agents and future AI technologies. I support this wide scope – it ensures that all forms of AI-driven tools and processes are in view, from internally developed models to third-party AI services. Notably, MAS explicitly excludes from scope any tools that operate on fixed, predefined logic (traditional deterministic programs or calculators). This clarity is welcome - it focuses the Guidelines on technologies with learning or inferential capabilities (which pose unique risks like unpredictability or bias) and avoids burdening simple software with AI compliance. In an era where AI capability is often embedded in off-the-shelf solutions, MAS's stance means FIs must account for vendor-provided AI tools and not just in-house models. Such breadth aligns with other leading jurisdictions' approach – for example, the EU's draft AI Act also casts a wide net over AI systems – and is wise to prevent loopholes where firms might otherwise outsource AI functionality to evade oversight. By defining model, system, and use-case distinctly, the Guidelines capture AI at different granularities. A "model" is an algorithm or method converting inputs to outputs (e.g. a machine learning model), while a "system" may comprise one or several models plus other components, and a "use case" refers to the specific application context. In practice, this means governance should consider not just the model in isolation, but the entire socio-technical system (data, software, hardware, and human process around it) and its real-world usage. I appreciate MAS's technology-agnostic stance – the Guidelines will apply across current and future AI tech, with MAS committing to update them as needed for new developments. This forward-looking approach is crucial given the pace of AI innovation. One suggestion is to ensure the scope explicitly covers AI components or external AI services that an FI might incorporate. For example, if an FI uses a pre-trained model from a third party or an AI-based data preprocessing service, these should fall under risk management scrutiny even if not full "systems" on their own. The broad definition likely already intends this (since third-party tools are mentioned), but clarity in examples would help FIs consistently identify such cases. Overall, MAS's proposed scope is suitably broad to encompass everything from credit scoring models to customer chatbots and fraud detection systems. It underscores that whenever a tool learns from data to generate decisions or content, the FI must treat it with the robust governance and ethics considerations that responsible AI demands. **Question 3:** I strongly endorse MAS's focus on Board and senior management accountability in AI risk oversight. In my three decades in governance and risk, one constant is that effective risk management starts at the top. The Board of Directors and C-suite must set a clear tone that the use of AI is not merely a technical experiment but a strategic initiative to be managed with the same rigor as credit, liquidity, or cyber risks. The proposed Guidelines rightly expect the Board and senior management to establish frameworks, policies and processes for AI risk management, and to foster an appropriate risk culture around AI use. This includes integrating AI considerations into the enterprise risk management framework – ensuring that existing risk policies (operational risk, IT risk, compliance, etc.) explicitly cover AI-related scenarios. In practice, the Board should approve an AI governance framework or policy that articulates the FI's risk appetite for AI (e.g. what applications are off-limits), roles and responsibilities, and oversight mechanisms.

		It is encouraging that MAS highlights Board and senior management must maintain effective oversight of AI risks and ensure clear accountability across the organization. To fulfill these duties, boards will need to invest in understanding AI – for instance, through training or bringing in advisors with AI expertise – so they can ask the right questions of management. A practical challenge is that AI, especially advanced techniques like deep learning or generative models, can be complex to comprehend. Thus, part of the Board's responsibility is ensuring they have access to expertise (perhaps via a technology advisory subcommittee or external experts) and that management reports on AI include not just technical metrics but plain-language explanations of risks, controls and incidents. Senior management, for its part, should designate clear owners for AI risk management (often a Chief Risk Officer or Chief Data/AI Officer working jointly) and ensure coordination between business units, IT, data science teams, and risk/compliance functions. I also advocate that leadership sets up internal communication and whistleblowing channels specific to AI issues. Employees should be empowered to escalate AI-related concerns – for example, if a data scientist detects bias in a model or an employee notices an AI tool being misused, they need safe avenues to report it. This supports a culture of transparency and ethical awareness, reinforcing accountability. In my experience, boards that encourage a "speak-up" culture – where staff can flag AI ethics or risk issues without fear – are far more effective at catching problems early (this ties into the broader whistleblowing frameworks that have served well in compliance areas). Finally, the Board should endeavor to regularly review AI risk reports and question whether AI use aligns with the institution's values and regulatory obligations. MAS's inclusion of Board oversight sends a strong signal: no AI deployment should proceed or continue without the knowledge and approval of the institution's highest governance bodies. This level of Board engagement, while demanding, is essential to ensure accountability and to uphold public trust – especially for FIs and other public-interest entities. Just as boards are accountable for financial soundness, they will be accountable for AI governance failures, so it is wise of MAS to set this expectation clearly. **Question 4:** The proposal to establish a dedicated cross-functional committee when an FI's overall AI risk exposure is deemed material is a sound governance measure. I support this idea: a specialized committee can provide focused oversight and draw in expertise from across the organization to holistically manage AI risks. Many AI risks are interdisciplinary – they touch on technology, ethics, customer impact, legal and operational issues – so no single department can effectively oversee all aspects. A cross-functional committee, composed of senior stakeholders from risk management, compliance, IT/cybersecurity, data science, business lines, and possibly legal or audit, ensures that diverse perspectives inform AI governance. This helps in integrating principles of fairness, transparency, human oversight, and data integrity directly into project reviews and risk assessments by having the relevant experts at one table. MAS asks how to assess if an FI's AI risk exposure is "material" enough to warrant such a committee. In practice, FIs should develop criteria as part of their AI risk framework. For example, the presence of any high-risk AI use cases (as determined by the risk materiality assessment), a significant number of AI systems in critical operations, or AI deployments that affect a large customer base could all signal material exposure. MAS's guidance (paragraph 3.10 of the Guidelines) on considering impact, complexity, and reliance across the portfolio will be useful. I would suggest that if any AI system is classified in the highest

		risk tier (due to high potential impact and autonomy), or if moderate-risk AI systems are pervasive in key business areas, the FI should err on the side of establishing the committee. In smaller organizations where a full new committee might be burdensome, the intent could be met by expanding the mandate of an existing senior management committee (e.g. an Operational Risk or Technology Risk Committee) to cover AI oversight explicitly. The key is not formality for its own sake, but ensuring dedicated attention and coordination at senior levels. Once in place, a cross-functional AI Risk Committee would typically be responsible for: defining the FI's AI strategy and risk appetite, approving or vetting major AI initiatives, monitoring the AI inventory and key risk indicators, reviewing incidents or near-misses, and ensuring remedial actions. MAS rightly notes that this committee should proactively address potential gaps in risk management. From my experience, such committees are most effective when they report up to the Board (or Board Risk Committee) regularly, thus creating a through-line of accountability. They should also set policies that business units must follow (e.g. requiring all AI projects to go through a risk assessment and be recorded in the inventory). One practical challenge is making sure this committee has sufficiently senior clout and expertise – it should ideally be chaired by a C-level executive (or even a Board member in a large FI) and include those with deep technical understanding of AI as well as those who understand customer and regulatory impacts. In sum, I concur that for FIs with significant AI use, a cross-functional oversight body is essential. It embodies the principle of accountability by designating a clear forum for AI risk deliberation. It also exemplifies human-in-the-loop governance - even as AI systems operate, humans (from diverse fields) collectively keep watch over them. MAS might consider providing further guidance or examples on what constitutes "material" AI risk exposure. I believe this principle to be well-placed. When implemented well, such committees will strengthen oversight and help ingrain a consistent, enterprise-wide approach to responsible AI use. **Question 5:** I fully support MAS's proposal that FIs establish clear definitions, criteria and processes – backed by robust systems – to identify all uses of AI in their business and functions. In my governance experience, you cannot manage what you haven't identified. Many organizations discover AI being used in corners of the business that senior management or risk teams were unaware of. A formal AI identification process ensures there is a structured way to surface every AI model, system, or tool in operation (or planned) across the enterprise. This should cover not only the obvious, large-scale AI applications (like an AI credit scoring engine), but also more modest uses (perhaps an AI plugin in a marketing tool, or an HR resume-screening AI). MAS's emphasis on consistent identification criteria and assigning clear responsibility for this task is prudent. Typically, FIs might assign liaisons in each business unit to report any AI usage, complemented by centralized oversight from an AI governance team. A practical challenge is the prevalence of "shadow AI" – analogous to shadow IT – where employees or departments use AI tools without formal approval. This could be something as simple as an analyst experimenting with a machine learning API, or using a generative AI service to draft code or reports, without informing risk management. Such unsanctioned usage can introduce significant risks (e.g. data leakage, uncontrolled decision-making) outside the purview of the FI's controls. Therefore, the identification process must include measures to detect and prevent shadow AI. Firms should update their internal policies to

		require that any new AI use (including third-party AI tools or cloud services) be declared and evaluated. Technical measures can help too: for instance, monitoring network traffic for calls to known AI service APIs, or restricting installation of unapproved software, analogous to DLP (data loss prevention) controls. MAS's Guidelines could explicitly mention shadow AI risk – indeed, identification is not just a passive inventory, but an active vigilance against unsanctioned AI deployment. Encouraging a culture where staff voluntarily disclose AI experiments is also key; this ties back to the risk culture set by the Board. If staff fear bureaucracy or reprimand, they may hide AI use – so FIs should instead create channels (even internal "AI sandboxes") where innovation is allowed but visible. To achieve consistent identification, clear definitions of what constitutes "AI usage" must be communicated enterprise-wide. MAS's definitions can serve as a base – for example, any application that learns from data or makes autonomous inferences qualifies. Firms might develop a simple questionnaire or decision tree for employees to flag AI (e.g. "Does this tool use machine learning or AI to generate outputs? If yes, register it."). Robust systems may involve an IT asset registry integrated with procurement: any purchase of software or subscription that has AI features triggers a notification to risk management. Periodic attestations from business heads can also help – requiring them to certify every quarter that all AI in their area has been reported. In summary, I agree with MAS that systematic identification is foundational. It upholds transparency internally, enabling all subsequent risk management steps. With clear criteria and strong processes, FIs can avoid the common pitfall of "unknown AI" lurking in the organization and ensure no AI system operates without the appropriate governance spotlight. **Question 6:** Building on identification, I believe that maintaining an accurate and up-to-date AI inventory is a cornerstone of effective AI risk management. I strongly support MAS's proposal for FIs to establish such an inventory. An AI inventory is essentially the centralized record of an institution's AI assets – it brings visibility and enables oversight. In practice, this inventory should list every AI use case, system, or model, along with key attributes for each. MAS rightly suggests including attributes like the AI's purpose, its risk materiality rating, and lifecycle status. I would add that it's valuable to record ownership (who is the accountable executive or team for that AI), the data it relies on, and links to relevant documentation (e.g. model documentation or validation reports). Over time, the inventory can serve as a living map of an FI's AI landscape – something board and management can review to understand where AI is deployed and concentrate oversight on higher-risk items. One practical recommendation is to use the inventory as more than a static list; it should be integrated into the AI governance workflow. For example, when a new AI project is proposed, adding it to the inventory (or a preliminary registry) should be one of the first steps, triggering a risk assessment process. Conversely, no AI system should go live or be procured without an entry in the inventory and appropriate approvals. This requires internal discipline but is achievable if mandated from the top. Many FIs already maintain model inventories for model risk management purposes; those can be expanded to cover AI systems broadly (including things that might not be traditional financial models, like an AI-driven customer chatbot). MAS even notes the inventory could be a standalone AI inventory or an enhancement of existing inventories, as long as clear linkages exist. The latter point is crucial – linking the AI inventory with other inventories (e.g. IT asset inventories, data inventories) can break down silos and ensure consistency. For instance,

		if there's a register of all software applications, any that include AI functionality should cross-reference the AI inventory. Maintaining the inventory "accurate and up-to-date" will require ongoing governance. FIs should establish a process for regular review (say, quarterly or semi-annually) of the inventory by the AI risk committee or a similar body, to verify that entries are current and complete. Each entry might be reviewed for changes: e.g. Has the model been updated or retired? Has its risk rating changed due to new information? Are new AI tools being used by departments that haven't yet been captured? In my experience, one challenge is ensuring completeness – something may slip under the radar. Mitigants include broad communication (so all staff know to report AI uses) and perhaps usage of discovery tools as mentioned earlier. Another challenge is the effort to populate and maintain details, but this is a necessary investment in governance. I also concur with the viewpoint that the AI inventory can double as a repository of risk management artefacts. For each AI system, it is helpful to attach its risk assessment results, validation reports, bias testing outcomes, and any approvals obtained. This way, the inventory is not just a list, but a auditable system-of-record for AI governance. It strengthens accountability and auditability – internal audit or regulators can look at the inventory and drill down into any listed AI system to find evidence of how it was vetted and how it's performing. Overall, the AI inventory is an essential tool for bringing the principles of transparency, data integrity, and oversight to life. MAS's guidance on this is very welcomed, and FIs should treat the inventory as the backbone of their AI governance program. **Question 7:** I understand that MAS proposes that FIs' AI risk materiality assessments consider at least three dimensions: impact, complexity, and reliance. I find this framework sensible as a minimum, and I appreciate MAS's focus on capturing different facets of AI risk beyond simple "high/low" notions. From a governance perspective, having a structured methodology to classify AI use cases by materiality is crucial to apply controls in a proportionate way. Let me address the following 3 specified dimensions:- 1) Impact This covers the potential consequences if the AI were to behave unexpectedly or fail. It includes impact on the FI (financial loss, operational disruption, reputational damage) and impact on customers or stakeholders (e.g. unfair outcomes, privacy intrusions). Impact essentially measures harm severity. I wholeheartedly agree this must be a primary dimension – it aligns with the principle of fairness and customer protection. A lending AI that could unfairly deny loans or an AI trading system that could cause large losses clearly have higher impact risk. 2) Complexity This captures how complex or opaque the AI system is, and how novel the technology or data is. Complexity often correlates with difficulty in oversight and understanding – for instance, a deep learning model on unstructured data might be considered high complexity because its inner workings are less explainable and it may behave in unexpected ways. Complexity can also include the difficulty of maintaining the model or the degree of changeability (adaptive models that evolve online are more complex to manage). Including complexity is wise because simpler models (like a straightforward regression) may pose less risk compared to highly complex algorithms, all else equal. This ties to explainability:

		the more complex, the harder to explain and thus the greater potential governance challenge. 3) Reliance This reflects the extent of the FI's dependence on the AI's decisions, i.e. the level of autonomy given to the AI and the criticality of its use. If an AI system operates with little human oversight (fully automated) or in a mission-critical capacity (e.g. directly executing trades or approving transactions), the reliance (and thus risk) is higher. Conversely, an AI that only provides a recommendation to a human, who can override it, implies less reliance. This dimension is fundamental to human oversight considerations: where reliance is high, we likely need stronger safeguards or fallback plans. Together, these 3 dimensions cover a lot of ground. However, MAS also inquires if other risk dimensions should be included. In my view, a few additional considerations could enrich the assessment: - Data Sensitivity & Privacy The nature of data an AI uses or produces could be an explicit dimension. For example, AI models processing highly sensitive personal data (think wealth management AI using personal financial data, or health-related data in insurance) carry privacy and legal risks, even if the model's impact or complexity is moderate. If not a standalone dimension, it should at least factor into impact (impact on customer privacy). Including it explicitly might encourage FIs to think about compliance with data protection laws as part of materiality. - Scale and Reach How widespread is the AI's use? An AI deployed across millions of customer accounts or across many business units presents a different risk profile than one used in a small pilot or a single department. Even a low-impact-per-instance AI can pose material risk if scaled massively (small errors multiplied many times). Thus, scale (number of users or volume of decisions affected) could be a dimension or a modifier in assessing impact. - Change Management / Controllability While perhaps more qualitative, one might consider how easily an AI's outcomes can be reversed or corrected if issues arise. For instance, if an AI decision is irreversible (e.g. automatically executes trades), that's higher risk than if outcomes can be reviewed and rolled back. This notion of reversibility or controllability could be part of reliance or separate. It asks: if the AI makes a mistake, how contained or correctable is the damage? MAS's current 3 dimensions are a solid starting point, and I would not want to overcomplicate the assessment with too many factors. Each FI could incorporate sub-factors tailored to their context. The Guidelines might suggest, for example, that FIs consider whether ethical and societal impact is relevant (some FIs might include a dimension for "ethical risk" if an AI could potentially cause discrimination or social harm beyond pure financial impact – though arguably that is part of impact as well). The key is ensuring a consistent, structured approach so that each AI use case is evaluated on these axes in a comparable way. It might also help if MAS provides illustrative criteria for low/medium/high in each dimension. Today, firms might interpret them differently: e.g., what qualifies as "high complexity"? Is it the algorithm type, the data, the lack of explainability? Some guidance or industry calibration (perhaps through examples in an Annex) would improve consistency across FIs. In summary, I concur that impact, complexity, reliance are necessary dimensions for assessing AI risk materiality, and I recommend considering data sensitivity and scale as additional factors. Adopting a multi-dimensional risk rating echoes the principle of proportionality – higher risk scores drive stronger controls. It also embeds explainability

		and fairness concerns into the risk metric (via complexity and impact). By implementing this, FIs can better prioritize oversight for AI systems that pose the greatest risk to consumers and the financial system, while allowing more flexibility for low-risk innovations. **Question 8:** I whole-heartedly support the holistic approach of the MAS's draft Guidelines' call for FIs to implement standards, processes, and controls across the entire AI life cycle, proportionate to risk. Managing AI risk isn't a one-off checkpoint; it must span from development through deployment to ongoing use and eventual decommissioning. The key areas MAS highlights – data management, fairness, transparency & explainability, human oversight, third-party risk, model/algorithm selection, evaluation & testing, technology and cybersecurity, reproducibility & auditability, review & monitoring, and change management – together cover the full spectrum of AI governance. This comprehensive list aligns well with internationally recognized AI governance principles (it echoes the FEAT principles of Fairness, Ethics, Accountability, Transparency) and frameworks like the NIST AI Risk Management Framework. **Question 9:** Overall, my personal view is that the draft Guidelines are thorough and well-crafted. I see them as an excellent foundation. However, I would like to highlight a few additional considerations and nuanced areas that, in my view, deserve attention in the final Guidelines or future guidance. These stem from an AI governance lens and practical challenges I've observed: 1) AI Ethics and Values Alignment - While the Guidelines focus on risk management mechanics, FIs should also ensure their use of AI aligns with their corporate values and ethical principles. This goes beyond avoiding bias – it's about the purpose for which AI is used. For instance, even a technically accurate AI could be deployed in ways that customers find invasive or that erode trust (such as overly aggressive personalized marketing). Boards should consider the ethical implications of AI use cases ("just because we can do it with AI, should we?"). Some firms establish Ethics Committees or seek input from independent ethics advisors on AI projects that may have societal impact. Incorporating a brief expectation in the Guidelines that FIs embed ethical deliberation in AI development (drawing on the "Ethics" in FEAT) would strengthen accountability. This could be part of risk culture: leadership signaling that AI must be used in service of customers' interests and societal good, not merely for profit or efficiency at any cost. 2) Whistleblowing and Speak-Up Mechanisms - As mentioned earlier, encouraging an internal culture where employees can raise AI-related concerns is vital. The Guidelines touch on risk culture, but a specific reference to whistleblowing frameworks could be beneficial. FIs might, for example, extend their existing whistleblowing or conduct hotlines to cover AI ethics and risk issues explicitly. If a data scientist feels pressured to deploy a model they believe is unfair, or an employee notices an AI making troubling decisions, they should know how to report this without fear of retaliation. Regulators in some jurisdictions (including the UK's Senior Manager Regime, etc.) have highlighted the importance of speaking up in risk management. MAS could encourage FIs to review their whistleblowing policies to ensure AI matters are covered. In practice, some firms have even instituted "AI ombudsman" roles or internal forums to discuss ethical dilemmas. The presence of

		internal checks and balances via empowered staff can significantly augment formal oversight structures. 3) Insider Threats and AI Misuse - A nuanced risk not explicitly spelled out is the insider threat regarding AI. This can take a few forms: a malicious insider might try to manipulate an AI system (e.g. feeding it biased data or altering outputs) to commit fraud, or could steal sensitive AI model information for personal gain or sabotage. Alternatively, well-meaning employees might misuse AI tools in ways that inadvertently violate policies (for instance, inputting confidential data into an external AI service, causing a data leak). FIs should incorporate AI systems into their insider risk programs – meaning monitor usage logs, restrict access rights, and have alerts for unusual activities involving AI systems (similar to privileged IT system monitoring). Training for staff on the proper use of AI tools (what data can/can't be input to public models, etc.) is important to prevent accidental insider-related incidents. While the Guidelines imply controls like access management and cybersecurity, making the insider dimension explicit would remind FIs to extend their insider threat controls to the AI arena as well. 4) AI Literacy and Training - One gap that could be emphasized is the need for building AI literacy across the organization. The Guidelines mention "capabilities and capacity" in passing (ensuring staff have competence for AI risk management). I would underscore that all three lines of defense – from development teams to risk/compliance staff to internal audit and up to the Board – need some level of AI understanding. This likely means FIs should invest in training programs about AI fundamentals, the firm's AI governance policies, ethical principles like fairness, and how to spot or handle AI incidents. A general awareness among employees will support effective implementation of these Guidelines. For instance, front-line staff who know the basics of AI risks are more likely to notice and report issues. The Board and senior management in particular should undergo targeted briefings so they can discharge their oversight duties knowledgeably. AI literacy can be seen as part of the risk culture – a workforce educated on AI's benefits and risks is a strong first line of defense. I recommend MAS encourage FIs to incorporate AI governance into their ongoing training curriculum. 5) Client/Customer Communication and Transparency - While the Guidelines largely focus on internal processes, responsible AI use also has an outward-facing aspect: transparency with customers or clients. I would encourage consideration of whether (for certain high-impact use cases) FIs should inform customers that an AI is involved in decision-making and provide avenues for recourse. For example, if an AI declines a customer's loan or detects a possibly fraudulent transaction, does the customer get an explanation or a human review option? Globally, there is movement towards model transparency to consumers (the EU AI Act will likely require some user notification and explanation for high-risk AI decisions). MAS's FEAT principles already emphasize transparency, so aligning the Guidelines with an expectation that outcomes of AI, when communicated to customers, should be clear and explainable is logical. This might be beyond the scope of these Guidelines which are prudential in nature, but it's worth noting as an emerging aspect of AI governance that FIs should prepare for. Ultimately, treating customers fairly in an AI-driven process includes giving them understanding and the ability to inquire or appeal decisions. 6) Continuous Monitoring of Emerging Risks - AI technology and threats evolve rapidly. The Guidelines could encourage FIs to establish an ongoing horizon-scanning function for new AI risks. For instance, today's concern might be bias in traditional models, but tomorrow's might be an AI agent that autonomously executes trades learning bad strategies, or new

		forms of algorithmic collusion. A suggestion is that FIs maintain an "AI risk register" or knowledge base that is periodically updated with identified risks and potential harms (somewhat alluded to by MAS's references to AI risks and harms in its papers). This living register ensures the institution's policies and controls adapt. It's an extension of the risk materiality exercise – not just classifying current uses, but anticipating future risk scenarios. Given the pace of AI, regulators and FIs will need a dynamic approach, and fostering that mindset early is beneficial. 7) Integration with Broader Frameworks - FIs should ensure that their AI risk management does not sit in isolation but is integrated with their overall governance. This means linking AI risk appetite with the enterprise risk appetite, incorporating AI oversight into existing governance structures (e.g. including AI in risk committee agendas, audit plans, etc.), and ensuring consistency with related regulations (data protection, outsourcing guidelines, model validation guidelines, etc.). Perhaps MAS could explicitly encourage FIs to cross-reference and integrate – for example, aligning AI model validation with existing Model Risk Management (MRM) guidelines, aligning AI cyber controls with MAS's Technology Risk Management Notice, and so on. The goal is a cohesive governance system where AI is another dimension of risk considered in all relevant management processes rather than a "silo" topic. In practice, this prevents gaps (like an AI system passing an AI ethics check but failing basic IT security, or vice versa). Given that many banks are likely to have mature risk frameworks, mapping the new AI requirements into those frameworks will be an important implementation step. **Question 10:** In my opinion, a one-year transition is reasonable but also ambitious, given the breadth of the said Guidelines. I believe it strikes a balance - providing time for FIs to adapt, without unduly delaying the uplift in risk management. Many FIs will likely need that time to conduct gap analyses, develop or update policies, and build up capabilities (such as inventories, new committees, training programs). Larger institutions may have a head start – some elements like model inventories or fairness checks might exist in parts – but smaller ones could be starting almost from scratch. From a governance standpoint, I would recommend that each FI's Board treat the issuance of the Guidelines as an urgent priority and task senior management with formulating a clear implementation plan. This plan should map each aspect of the Guidelines to actions, responsible owners, and deadlines within that 12-month window. Boards should oversee this plan closely, receiving progress reports (perhaps quarterly) to ensure the FI stays on track. One year can pass quickly, so FIs might need to prioritize. I would think it's advisable to focus on establishing the foundational pieces early. For example, identify all current AI uses (to know your exposure), appoint accountable officers or committees, and implement any high-impact controls for existing AI systems as soon as possible. It might be pragmatic to take a phased approach within the 12 months – addressing the highest-risk AI applications with new controls first, and then broadening out. MAS could consider issuing further guidance or hosting industry workshops during the transition period to assist, especially for smaller FIs that might benefit from templates or examples (for instance, examples of AI inventory structures or risk assessment frameworks). It's worth noting that internationally, regulatory expectations on AI are also ramping up, so this 12-month timeline aligns well with global momentum. The EU's AI Act, for instance, is on the horizon; FIs operating globally will likely be aligning with multiple regimes.

		Starting compliance efforts now, under MAS's timeline, will position Singapore FIs ahead of the curve in AI governance. If there is any concern, it might be whether 12 months is enough for FIs with very low maturity in AI governance. However, since the Guidelines are risk-based, an FI with minimal AI usage will have a lighter lift (mainly putting basic policies in place – which is doable in that time). An FI with heavy AI usage, conversely, likely already recognizes the need for such controls and may have some elements to build on. On balance, I believe 12 months is fair and in line with similar transitions (many new regulations allow about a year). It's also reassuring that MAS plans to actively engage and monitor – it is likely that we can expect supervisory scrutiny on AI risk management once the period lapses. This external deadline and oversight will likely motivate FIs to allocate appropriate resources and priority.
28	Endowus Singapore Pte Ltd	**Question 1:** We agree that rules should fit the risk level, and acknowledge the guidance and examples provided in the Annex regarding proportionate application. To further assist FIs in classifying use cases that may not fit neatly into the provided examples, we request that MAS explicitly confirms that effective 'Human-in-the-Loop' review (where human validates the output before use) serves as a primary determinant for classifying a use case as 'Low Risk' or 'Not Integrated'. We want to ensure that standard productivity tools used by our staff are not treated with the same heaviness as critical trading algorithms, but are still subjected to the appropriate security measures. This would provide FIs with a clear principle to apply to novel or evolving AI tools not yet listed in the Annex. **Question 2:** We note and welcome the clarification in Paragraph 1.2(d) that tools based on predefined logic are excluded. To further avoid ambiguity for the Wealth Management sector, we suggest explicitly listing 'standard, rule-based investment formulas and portfolio construction methods' as examples of this exclusion. This ensures that fixed mathematical rules are clearly distinguished from Machine Learning models. **Question 3:** We support Board oversight. However, for FinTechs and smaller FIs, assigning specific AI responsibilities to the Board requires flexibility. We suggest that the Board's role should focus on Risk Appetite and Strategy, rather than technical lifecycle approval. We recommend to allow Senior Management to delegate the 'approval of frameworks' to a specialised Risk Committee rather than requiring full Board approval for technical policy updates, to ensure agility in a fast-moving tech environment, while still keeping the Board members in the loop. **Question 4:** There should be caution regarding the mandate for a separate dedicated committee. For agile FIs, AI risk is best managed within existing Technology Risk or Enterprise Risk Committees to avoid siloed governance. Similar to the recommendation above, we propose that FIs be permitted to expand the mandate of existing Risk Committee to cover 'Material AI Risk', rather than being forced to form a new body. Regarding 'Organisational

		Level Assessment', this should be defined by the impact on client outcomes (e.g., percentage of AUM managed by AI) rather than just the technical complexity of the model. **Question 5:** 1. Regarding Paragraph 3.3, we request flexibility for FIs to assign 'control function' responsibilities to existing departments (e.g., Technology Risk or Compliance) rather than requiring the designation of a new, standalone AI control function. This aligns with a proportionate approach for smaller FIs. 2. Identifying 'Embedded AI' in third-party SaaS tools (e.g., CRM software, Slack, Office 365) is challenging as vendors increasingly add AI features silently. We suggest MAS distinguish between 'Proprietary/Core AI' (models we build or control that impact investment decisions) and 'Incidental/Embedded AI' (AI features within standard SaaS tools). The identification requirements for Incidental AI should be lighter, relying on vendor certifications rather than internal 'robust systems'. **Question 6:** We support the inventory requirement. However, maintaining a real-time inventory for rapid-cycle development (Agile/DevOps) is resource-intensive. We request that the inventory update frequency be tied to deployment cycles or material changes rather than requiring continuous real-time updates. We also ask for flexibility to integrate this into existing Software Assessment Management (SAM) register rather than creating a duplicate 'AI Registry'. **Question 7:** We agree with the proposed risk dimensions (Impact, Complexity, Reliance). We support the approach in Paragraph 3.9 to assess 'residual risk'. This allows FIs to account for specific mitigating controls, such as the ability to 'undo' or reverse a transaction, or 'human-in-the-loop' interventions, which naturally reduce the risk profile. Relying on residual risk assessment is more effective than adding 'Reversibility' as a standalone dimension. **Question 8:** Regarding Third-Party AI, FIs often lack the leverage to demand detailed model weights or training data transparency from global companies (e.g., Microsoft/OpenAI/Google). We ask MAS to accept "Third-Party Audits/Certifications" (e.g., ISO 42001, ISO 27001, SOC2) as sufficient due diligence for Tier-1 General Purpose AI providers, as obtaining 'white box' transparency is often commercially impossible. Regarding the term 'fairness', in wealth management, 'bias' is different from lending. We differentiate clients based on risk profile (which is intentional). Please clarify that 'fairness' controls relate to protected characteristics (race, gender) and not financial profiling required for suitability assessments. **Question 9:** 1. Third-Party Disclosures & Competition: The Guidelines should acknowledge that FIs cannot compel major global vendors (e.g., OpenAI, Microsoft) to disclose proprietary model weights or training data. If strict 'white box' transparency is enforced, only the largest incumbent banks with significant bargaining power will be able to secure such

		custom disclosures. This put smaller FIs and FinTechs at a severe competitive disadvantage, effectively barring them from leveraging AI at scale. We strongly urge MAS to accept public certifications (e.g., ISO, SOC 2) as sufficient due diligence for Tier-1 General Purpose AI. 2. Shadow AI & Scope of 'Personnel': We note that Paragraph 5.1 focuses on ensuring the competence and conduct of personnel 'involved in developing and deploying' AI. However, 'Shadow AI' risks often arise from general staff using public AI tools on personal devices, who do not fall under the 'developer/deployer' definition. 3. We request that MAS explicitly clarifies the expected control standards for this specific risk. Specially, we seek confirmation that FIs may rely on conduct-based controls (e.g., Acceptable Use Policies, training, and staff attestations) to manage Shadow AI risks on personal devices/accounts, rather than being expected to implement technical surveillance or blocking measures on non-corporate devices, which present significant feasibility and privacy challenges. **Question 10:** While we aim to comply with the proposed 12-month transition period, we wish to highlight that the retrospective application of the new guidelines to established legacy models and systems may require significant lead time, particularly where dependencies on third-party provider roadmaps are involved. We would welcome supervisory guidance or clarification in the response to this consultation that reinforces a risk-based implementation approach during the transition. Specifically, we suggest confirming that FIs should prioritise resources towards High Materiality and Critical AI uses during the initial transition phase, ensuring that the most significant risks are addressed first.
29	Ensign InfoSecurity	**Question 1:** It is appropriate to consider application of the guidelines to all FIs in a proportionate manner. Points for consideration include: 1) Criticality of the business process in which the AI system is employed/deployed for. 2) Access to sensitive information, e.g., personal data, financial transactions, contracts. 3) Differentiation between decision-support and decision execution. 4) Level of human intervention, e.g., human-in-the-loop, human-on-the-loop, human-out-of-the-loop. 5) Scale and magnitude of impact to stakeholders in failure or breach of AI system for the relevant business processes. 6) AI system used for experimentation or proof-of-concept vs. production. 7) Implication and / or dependency towards external stakeholders, ecosystem partners and the wider supply chain. The applicability of the AIRG, as illustrated in Figure 2, could consider the following, as there would be necessarily different risk responses and management approaches: 1) The AI system is used as a commercial-of-the-shelf (COTS) solution; 2) The AI system is adapted/customised on top of a COTS solition; 3) The AI system is internally developed.

		Question 2: The scope is appropriate and aligned with broader industry perspectives. It may be useful to clarify if the AIRG applies to experiments and proof-of-concepts which are applied on real-world contexts (e.g., scenarios and processes), but may not interact with production or real-world data **Question 3:** The proposed responsibilities of the Board and senior management are consistent with the MAS Technology Risk Management Guidelines, and has taken into consideration of the multi-disciplinary and AI system lifecycle nuances. There is room to consider highlighting the following: 1) The scale and speed of access to data AI systems may have, and its implications to risks, in aggregate or in "toxic" combinations. 2) The sensitivity of data AI systems have access to or create. The above points suggest the need to highlight aggravated risk exposure, especially with agentic AI systems (which embodies automation and insights). There may be a need to highlight that the Board and management consider the following as well: 1) Cost of AI system dependency failure and its alternative (e.g., substitution or business continuity management) 2) Concentration risks across the organisation 3) Level of control and / or influence to the employment / deployment of the AI system (e.g., COTS or SaaS vs. inhouse developed solution). **Question 4:** **Question 5:** This is a necessary guideline to provide to FIs. Many FIs already perform robust business impact and risk assessments to manage both operational and non-operational risks. It may be useful to suggest the following: 1) Leverage business impact assessments (BIAs), which are typically performed regularly, e.g., annually, to include identification of business processes which depend on AI systems, and document the AI systems used. 2) Suggest a AI use case inventory template. 3) Assess / reassess materiality of business processes to AI systems dependency. 4) Document risk appetite (at the functional, business unit, and organisation levels) 5) Establish risk responses to address identified risks The above should be aligned to ERM practices that FIs already practice, e.g., COSO and COBIT. **Question 6:** **Question 7:**

		The suggested risk dimensions are appropriate. However, due to the multi-disciplinary nature of AI and the wide applicability in AI systems use, it is relevant to consider additional dimensions offered in robust enterprise risk management (ERM) frameworks to avoid constraining AI risk management to be seen as purely technology-related risk management, but also a business risk management aspect. Taking this approach will allow for a more well-rounded appreciation of risk exposure and risk management by also complementing it with additional lenses, also suggested by MAS' FEAT principles, Outsourcing Guidelines, in addition to Technology Risk Management Guidelines. **Question 8:** Notable references include: 1) NIST AI RMF - NIST AI 100-1 (https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf) 2) ISO/IEC 23894:2023 (https://www.iso.org/standard/77304.html) 3) ISO/IEC 42001:2023 (https://www.iso.org/standard/42001) **Question 9:** Generally, it may be useful for MAS to consider framing AI-related risks as a constituent business risk rather than narrowly constraining it as a technology risk. This is mainly due to the multi-disciplinary and wide applicability nature of the technology and science which can be applied broadly across an organisation. **Question 10:** The proposed transition period of 12 months is minimally necessary, but may be subjective, depending on the level of integration a FI may have employed / deployed AI system(s) in the organisation. Some language for open consultation with MAS on specific challenges in meeting the transition period should be included.
30	Ernst & Young	**Question 1:** EY supports MAS's proportionate approach in AI Risk Management based on the key principles of Fairness, Ethics, Accountability and Transparency (FEAT). [Annex, page 29] On governance in paragraph 5(a), we suggest that a senior management governance body or committee be formed to oversee the development and adoption of AI to strengthen accountability. For the purpose of practicality and clarity, we suggest framing paragraph 5(b) as risk-based: rather than a blanket requirement for human review and validation of all AI outputs. the Annex should state that FIs set out the extent of requirements for human review and validation commensurate with the risk, materiality and context of each use case, with appropriate pre-deployment and periodic validation. Further clarity should be provided on whether the term “use” in paragraph 5(b) refers to go-live/implementation of an AI tool or each output. One of the key AI risks arise from use of unapproved AI tools. In paragraph 5(b), the Annex should include expectations on preventing and detecting unapproved AI usage.

		To avoid confusion with MAS's AI Risk Management Guidelines (AIRG) in Section 5b, 5c, 5d and 5e, we propose using "internal AI policy", "standards", "reference" or "leading practices") instead of "guidelines" within this Annex. **Question 2:** [Section 1, page 12 & 13] The definition of AI model may be misleading to the extent that the output of a model is neither a decision or recommendation, but rather a mathematical inference or prediction. Whether or not the predicted/inferred output is treated as a decision, recommendation or estimate depends on the way the AI model is used by the system it is integrated into. EY would suggest aligning with known global definitions such as IS/IEC 22989:2022. Under the definition of Trained model (5.11.7), "[...] A machine learning model is a mathematical construct that generates an inference, or prediction, based on input data or information. The trained model should be usable by an AI system to make predictions based on production data from the area of interest. [...]". **Question 3:** EY's own AI framework also emphasizes the responsibilities of board and senior management in overseeing AI risk management. As such we are supportive of the MAS aims and have no further comments. **Question 4:** **Question 5:** [Section 3.2, page 17] EY would suggest that, wherever appropriate, it should be clearly documented how the definitions and criteria used by the FI align with common standards or regulatory specifications. **Question 6:** [Section 3.4/3.5, page 17] EY welcomes the MAS guidance to require FI to maintain an AI inventory and agree that it is a foundational control for effective governance and oversight. To enhance clarity and usability in this area, MAS could consider the following: o Providing a clearer list of core attributes with optional enhanced attributes that firms can customize to fit their needs, ensuring a consistent baseline while maintaining proportionality. o Clarifying that AI inventories should be regularly updated through change management processes and checked against production systems, making it a living document rather than a one-time task.

		o Providing additional examples for "key roles and responsibilities" and encouraging alignment with existing three-lines-of-defence and senior management accountability frameworks, to strengthen governance oversight. These refinements would help ensure that AI inventories are both proportionate and operationally useful, while supporting MAS's objective of achieving robust, risk-based governance across diverse AI use cases in the financial sector. **Question 7:** [Section 3] The AI risk management approach and controls should be aligned to the risk level of the use case in Section 3.8. The MAS risk dimensions of impact, complexity and reliance form a solid baseline but do not explicitly assess whether an AI system's outputs are sufficiently verifiable to support effective control design and oversight. MAS may like to consider adding Uncertainty as a dimension, defined as the extent to which the AI application's outputs can be independently validated, measured with precision, and reproduced or corroborated using reliable alternative evidence. This dimension addresses scenarios where outputs are hard to verify even when impact and reliance appear moderate (e.g., human-in-the-loop, non-final decisioning), creating a hidden risk that warrants stronger mitigating controls Further, we would like to highlight that the complexity dimension is dependent on the organisation maturity, and that may not adequately reflect the materiality and the level of risk FIs may face. **Question 8:** Section 4 - Evaluation and Testing: Given that the document already includes specific AI use case examples, it would be helpful if it could also provide reference evaluation measures together with indicative performance thresholds, at least for the example use cases cited. These can be framed as non-exhaustive reference tests and illustrative thresholds. This would give FIs a clearer understanding of regulatory expectations. FIs are expected to detect data, concept, and model drift (see 4.23, model drift is unique to the new risks GenAI brings, but isn't covered in 4.14). Additional guidance on model drift methods and measurements (such as reproducible behavioral drift checks, including golden (ground truth) datasets or equivalent fixed evaluation anchors) would be helpful. As GenAI models are known to be unpredictable, clear guidance on monitoring requirements would help FIs in attaining a sufficient level of divergence detection in model behaviour. Section 4 - Reproducibility and Auditability: It may be challenging for Gen AI models to be fully reproducible due to the probabilitis nature, and the pushing of version updates and releases. A risk-based approach is suggested.

		For Auditability, additional guidance should also be provided on log management and AI evaluation. Immutable logging and data provenance monitoring are examples of risk controls for consideration. **Question 9:** EY has additional comments on the other aspects of the Guidelines: 1. Developing AI core skills and competences are central to the effective risk management of AI. Additional guidance should be provided in this area. For example, AI reviewers should include personnel with sufficient AI technical specialisation to understand and challenge architectural, modelling and system-level modifications. 2. LLM may not have explainability metrics as opposed to the case for Machine Learning (e.g, SHAP and Lime), so meeting this requirement might be very challenging. 3. EY cautions on having complexity as part of the risk assessment, as it does not necessarily reflect the materiality and level of risk of a given application. Also complexity is dependent on the maturity of the organisation in the adoption of AI. **Question 10:** EY is supportive of MAS proposal that a transition period of 12 months be provided and have no further comments.
31	Fibish Madathil	**Question 1:** We support MAS' proposal to apply the Guidelines to all FIs in a proportionate, risk-based manner, as set out in paragraph 1.5 and the Annex. The distinction between AI used as an integrated part of business processes and AI used in an assistive manner is clear and practical, and the examples provided are helpful in guiding consistent application. Requiring all FIs to implement basic AI policies, while calibrating governance, lifecycle controls, and capabilities based on the materiality of AI risks, appropriately balances risk management with responsible innovation. To further support consistency, MAS may consider providing additional non-prescriptive indicators, such as AI use in regulated activities, degree of autonomy, and impact on core business processes. **Question 2:** We support the proposed scope of AI use cases, systems and models for the application of the Guidelines. The definitions of AI models, systems and use cases are clear, technology-neutral, and sufficiently broad to cover current and emerging AI technologies, including Generative AI and AI agents. The exclusion of deterministic, rules-based tools is appropriate and avoids over-regulation of conventional systems. Overall, the proposed scope provides clarity while remaining flexible to accommodate future AI developments. **Question 3:** We support the proposed responsibilities of the Board and senior management in overseeing AI risk management. The expectations appropriately emphasise accountability, effective oversight, and integration of AI risks into existing enterprise risk management frameworks. The requirement for Boards to maintain sufficient understanding of AI to provide effective challenge is appropriate, as is the focus on fostering a sound risk culture. Overall, the proposed responsibilities are clear, proportionate, and aligned with good governance practices.

		Question 4: We support the proposal for FIs to establish a dedicated cross-functional committee where overall AI risk exposure is deemed material. Such a committee can enhance coordination across business, risk, technology, legal and compliance functions, particularly for complex AI use cases such as Generative AI and AI agents. To support consistent application, MAS may consider providing non-prescriptive guidance on assessing overall AI risk exposure at the organisational level, for example based on the concentration of high-risk AI use cases, use of AI in critical or regulated activities, degree of autonomy and reliance on AI outputs, and aggregate customer or financial impact. **Question 5:** We support the proposal for FIs to establish clear definitions, criteria and processes, supported by robust systems, to ensure consistent identification of AI across the organisation. Effective AI identification is foundational to governance, risk assessment and lifecycle controls, and helps mitigate risks from unapproved or unmanaged AI use. **Question 6:** We support the proposal for FIs to establish and maintain an accurate and up-to-date inventory of all AI usage. An enterprise-wide AI inventory is essential for effective governance, risk oversight, and lifecycle management, and enables aggregation of AI risks across the organisation. The flexibility to integrate AI into existing inventories is appropriate and pragmatic. **Question 7:** We support the proposed risk dimensions of impact, complexity and reliance as appropriate core factors for AI risk materiality assessments. These dimensions capture key considerations relating to potential harm, technical uncertainty and degree of dependence on AI outputs. In addition, where relevant, FIs may consider incorporating complementary dimensions such as data sensitivity, degree of autonomy, and third-party or concentration risk to reflect risks arising from sensitive data use, AI agents, and reliance on external AI providers. **Question 8:** We support the proposed standards, processes and controls across the entire AI life cycle, and the expectation for FIs to apply them in a proportionate manner based on risk materiality. The coverage of data management, fairness, transparency and explainability, human oversight, third-party risk management, evaluation and testing, monitoring, and change management is comprehensive and aligned with good practice. The proportional, risk-based application appropriately balances effective risk management with operational practicality and innovation **Question 9:** The Guidelines are comprehensive and well-structured. As additional considerations, MAS may wish to provide further guidance on accountability where human decision-makers override or rely on AI outputs, governance of continuously learning or dynamically updated AI systems, and enterprise-level aggregation and reporting of AI risks to support Board oversight. Clarification in these areas would further strengthen effective and consistent implementation.

		Question 10: We support the proposed 12-month transition period, which provides FIs with a reasonable timeframe to assess AI usage, establish inventories and governance frameworks, and implement proportionate lifecycle controls. A risk-based, phased approach within this period would be appropriate, particularly for larger or more complex institutions.
32	FIL Investment Management (Singapore) Limited	**Question 1:** Nil **Question 2:** Nil **Question 3:** We request clarification from MAS on what would be considered an acceptable "AI risk appetite" structure (e.g., mandatory categories such as model performance, conduct/fairness, security, privacy, third-party concentration, and GenAI-specific risks). We would also appreciate guidance on whether MAS expects AI risk appetite thresholds to be set at the use case level, per model family or platform, and/or at the overall portfolio level. **Question 4:** Nil **Question 5:** Nil **Question 6:** We note the use of the terms "Comprehensive" and "All AI use cases" in the Guidelines. With AI now embedded across many commonly used applications (e.g., Excel, Zoom, Workday), maintaining an inventory of every minor AI feature would create significant administrative burden and quickly become outdated. As a comprehensive inventory for embedded AI is impractical, we propose replacing it with a tiered inventory or category-level registration framework, supported by appropriate guardrails. **Question 7:** Nil **Question 8:** We seek clarification from MAS on how FIs can gain adequate assurance when using off-the-shelf AI solutions, specifically regarding training data rights (e.g., IP ownership, licensing), training data provenance (particularly where vendors cannot or will not provide disclosures at a level that enables integrity verification) and suppliers' adherence to fairness and security requirements. We also refer to footnote (14) and would appreciate clarification on the minimum controls expected before a pilot or partial deployment may proceed (e.g., baseline testing, threat assessment, data controls, human oversight). It would also be helpful for MAS to outline

		the conditions that would trigger the requirement for full pre-deployment independent validation. Additionally, we request clarification on MAS's expectations regarding kill switches, noting that a true kill switch may be infeasible for certain architectures. For high-risk AI, it would be helpful for FIs to have clearer guidance on the circumstances under which MAS expects a true kill switch (i.e., immediate system disablement), as opposed to alternative controls (such as rate limiting, fallback routing, feature flags, or scoped disablement). We note that the Guidelines require FIs to ensure rigorous validation regardless of whether AI models are developed in-house or procured. However, demonstrating conceptual soundness for black-box models is often not technically feasible, and full right to audit access is typically impractical for off-the-shelf solutions. We therefore propose replacing these requirements with provider attestations or independent third-party assurance, supplemented by last-mile controls implemented by FIs. Finally, for AI agents, we seek clarity on the minimum acceptable monitoring coverage across the action space (e.g., API calls, database queries, payments, messages), and how such monitoring should be scoped in accordance with proportionality principles. **Question 9:** Nil **Question 10:** Nil
33	FinbotsAI	**Question 1:** I support the proportionate, risk-based application of the Guidelines as set out in paragraph 1.5 and the Annex. This approach is necessary given the diversity in size, complexity, and AI adoption maturity across financial institutions. In practice, a key challenge is not proportionality itself, but its operational application over time. AI initiatives often begin as pilots or assistive tools and later transition into production workflows without a corresponding reassessment of risk materiality or escalation of controls. This silent transition - rather than initial experimentation - is where material risk frequently emerges. It would therefore be helpful for the Guidelines to explicitly encourage institutions to reassess proportionality when AI use cases evolve in scope, criticality, or decision impact, rather than treating proportionality as a one-time determination at initial deployment. In practice, institutions often observe that AI initiatives initially approved as pilots or decision-support tools gradually become embedded in operational workflows. This typically occurs through incremental scope expansion rather than formal re-approval and is where risk controls most often lag actual usage. **Question 2:**

		The proposed scope of AI use cases, systems and models is appropriate and sufficiently broad. It reflects how AI is increasingly embedded not only in standalone models, but also in end-to-end systems and decisioning workflows. One area that may benefit from further clarification is the distinction between: * exploratory or pilot AI use cases, * assistive AI (including Generative AI used to support human decision-making), and * AI systems embedded in regulated or customer-impacting decisions. While all three fall within the scope of AI risk management, they present materially different risk profiles. Explicit recognition of these differences would help institutions apply the Guidelines more consistently and proportionately. **Question 3:** I support the proposed responsibilities of the Board and senior management in overseeing AI risk management. Clear accountability at the highest level is essential given the potential impact of AI on customers, financial outcomes, and institutional trust. In practice, effective oversight requires more than formal approval or periodic reporting. As AI systems evolve post-deployment, Boards and senior management benefit from assurance mechanisms that provide ongoing visibility into performance, behavioral drift, and emerging dependencies, including reliance on third-party or foundation models. The Guidelines may therefore wish to emphasize continuous, evidence-based oversight as a complement to pre-deployment approval, without expanding the formal scope of Board responsibilities. In practice, Boards often receive assurance at the point of deployment but have limited visibility into how AI behavior evolves over time, particularly where models are retrained, externally updated, or repurposed for adjacent use cases without formal escalation. **Question 4:** The proposal to establish a dedicated cross-functional committee where overall AI risk exposure is material is sensible. AI risk typically cuts across technology, business, risk, compliance, legal, and data domains, and cannot be effectively managed within a single function. At the organisational level, overall AI risk exposure could be assessed based on a combination of: * the number and criticality of AI systems in production, * the extent to which AI influences regulated or customer-impacting decisions,

		* reliance on third-party or externally governed AI components, and * the institution's ability to monitor, explain, and intervene in AI-driven outcomes. This allows flexibility while providing a structured basis for determining when enhanced governance is warranted. **Question 5:** I support the proposal for financial institutions to establish clear definitions, criteria, and processes to consistently identify AI across all relevant business and functional areas. In practice, challenges often arise where AI capabilities are embedded within broader systems, vendor platforms, or decisioning workflows, rather than existing as clearly labelled "AI projects". Without clear identification criteria, such usage can be overlooked. Robust identification processes, supported by systems and internal guidance, are therefore critical to ensure comprehensive visibility and effective downstream risk management. **Question 6:** The proposal to establish and maintain an accurate and up-to-date inventory of AI usage is appropriate and necessary. An AI inventory is foundational to effective governance, risk assessment, and oversight. To remain effective, such inventories should be treated as living artefacts rather than static registers. In particular, inventories should be updated when: * AI use cases change in scope or decision impact, * models or systems are materially updated, or * dependencies on third-party or underlying models evolve. This helps ensure that inventory accuracy keeps pace with real operational risk. **Question 7:** The proposed risk dimensions of impact, complexity, and reliance provide a strong foundation for AI materiality assessments. Based on recent deployment experience, it may be useful to explicitly recognize additional dimensions that materially influence risk, including: * dependency on underlying foundation models that are governed outside the institution, * data boundary and data sovereignty considerations, * reversibility and fallback viability, and

		* economic and infrastructure sustainability for AI systems deployed at scale. These factors often determine whether AI risk can be managed effectively in production environments. In practice, reliance risk often manifests not from the number of AI systems deployed, but from the concentration of decision influence within a small number of models or platforms, particularly where multiple workflows depend on a common underlying model or vendor capability. **Question 8:** I support the application of standards, processes, and controls across the entire AI life cycle, applied in a proportionate manner. In practice, lifecycle risk often intensifies after deployment, as AI systems encounter real-world data, operational pressures, and evolving dependencies. Ongoing monitoring, recalibration, and performance validation are therefore as important as design-time controls. The Guidelines appropriately emphasise lifecycle coverage. Explicit recognition of post-deployment assurance and intervention mechanisms would further strengthen their practical effectiveness. From an implementation perspective, issues such as data drift, changing customer behavior, and operational constraints often surface only after systems are live, reinforcing the importance of continuous monitoring rather than reliance on initial validation alone. **Question 9:** One area that may warrant explicit recognition is the risk arising from dependency on underlying foundation models, particularly large language models, that are developed, updated, and governed outside the financial institution. Changes at this layer can propagate behavioural or operational risk into downstream systems, even where application-level controls remain unchanged. Traditional third-party risk frameworks may not fully capture this inherited and transitive risk. Explicit acknowledgement of this risk would help institutions extend existing controls without fundamentally altering the Guidelines' structure or intent. In practice, institutions are beginning to observe that changes in underlying models - such as version updates or modified safety behavior - can alter downstream system performance without corresponding changes in application logic, creating inherited risks that are difficult to detect through traditional validation approaches. **Question 10:**

		The proposed transition period of 12 months appears reasonable for most institutions, particularly where foundational governance, inventory, and lifecycle processes need to be established or enhanced. In practice, implementation timelines may vary depending on the complexity and scale of existing AI usage. A risk-based supervisory approach during the transition period would allow institutions to prioritise high-impact use cases while progressing toward full alignment. Overall, the proposed transition period strikes an appropriate balance between urgency and practical feasibility.
34	FIX Trading Community	**Question 1:** We think this is a reasonable approach and support MAS's proportionality intent. We recommend explicitly anchoring the Guidelines to established, operationally proven constructs already used by FIs for market integrity and operational resilience and urge consistency of definitions and implementation across different regulatory jurisdictions. In particular, we believe the proportionality logic should map cleanly onto: • MiFID RTS 6-style governance for algorithmic systems (controls, testing, change management, oversight) and • DORA-style ICT (information and communication technology) and operational resilience (including third-party dependencies and criticality-based scaling). We would also like to note that footnote 6 on page 13 brings in scope all regulated activities. This would include trading/execution activities and suggest that this be included in the footnote for clarity. It would also be helpful to make clear the responsibility of asset managers using a broker's AI systems, including where that relates to trading. **Question 2:** We are broadly supportive of the scope and definition statement here. We would like to note, however, that some of out of scope 'calculators' can introduce similar problems to AI as defined so there are aspects of the guidelines that may be worth considering. We also note that calculators may make use of data learnt by an AI system 'offline' (i.e., separate to the operation of the calculator) and could be used to influence the behaviour of the calculator. A theoretical example of LLM data usage is where LLM-derived signals are incorporated into an order scheduling algorithm. For example, the system could: • Aggregate and summarise relevant news for a given instrument prior to market open. • Derive a single sentiment indicator (e.g., positive/negative/neutral). • Adjust the intraday execution profile accordingly: if sentiment is strongly positive or negative, front-load execution earlier in the session and reduce participation later; if sentiment is neutral, follow the established historical schedule. • Consume the sentiment indicator once at algorithm start-up, rather than re-evaluating continuously intraday. Issues and considerations are as follows: Although the execution logic remains deterministic, the upstream input may vary between runs. LLM-generated outputs can be non-deterministic, which may lead to different trading patterns for the same instrument

		under similar conditions. This creates additional requirements for governance and control, including robust logging and auditability of: • The raw source content used (articles/headlines, timestamps, and providers). • The LLM prompt, model/version, and configuration parameters. • The generated summary and sentiment output provided to the algorithm. • The resulting schedule decision and any overrides or fallbacks applied. If the LLM is continuously re-evaluated intraday and its outputs feed directly into market access system workflows, its role shifts from a one-time pre-trade input to an active driver of execution behaviour. This increases complexity and risk, as execution decisions may change in response to evolving information and short-term noise, reducing predictability and making outcomes harder to test in advance. It also tightens the coupling between the LLM and the execution layer, requiring stronger real-time controls and oversight, and is more likely to be viewed as part of the trading decision process itself rather than a supporting analytics tool. Consideration should also be given to system design, such as catering for the failure of any AI components and their knock-on effects to other systems. Further to the above, we feel that industry-agreed definitions and a stable taxonomy are required, particularly when describing the usage of AI systems. If firms are allowed to define "AI" for themselves, inconsistency and regulatory arbitrage will be a more likely outcome. Our recommendations are: • MAS should reference or adopt a globally recognised AI taxonomy (or require an industry-agreed one) that covers model/system/use-case distinctions, levels of autonomy and the definition of terms such as "agentic AI". We note in particular that a taxonomy for AI usage (e.g., list of business processes) would be helpful and we refer to our response to Q5 for more on this point. • MAS should encourage reuse of established "trust signals" rather than bespoke one-off regulations. For example, provenance standards such as the Coalition for Content Provenance and Authenticity (C2PA) are relevant as a pattern for data and content authenticity disclosures (even if adapted for financial contexts) – (https://c2pa.org/) Firms cannot define AI for themselves; doing so creates inconsistency and regulatory arbitrage. • Agentic AI and autonomy levels should be better defined; our suggestion is to look into CSA (Singapore) consultation on Agentic AI - https://isomer-user-content.by.gov.sg/36/703ff9fe-9db1-4e09-98c2-89e3d7007ef0/Draft%20Addendum%20on%20Securing%20Agentic%20AI%20[For%20Public%20Consultation].pdf. Having very broad guidelines runs the risk of FIs implementing the guidelines in very different ways, with some 'racing to the bottom' and doing very little, and others 'racing to the top' and incurring a significant regulatory implementation burden. Using the example of trading, this clearly being an important topic, one subject to a great deal of technological innovation and global in nature, we believe that more specific guidance would be beneficial, ideally developed in collaboration with industry participants and the global regulatory community. We would be happy to support any such efforts. We note the publication is ISO 24030 on AI use cases as a useful reference.

		Question 3: We think this is a reasonable approach and is similar to the approach taken for managing, for example, algorithmic trading in other jurisdictions. We have two comments to make: • We note that responsibilities for communication to staff (as per annex) aren't clearly defined here • We support MAS's direction but recommend strengthening the text so boards are accountable for governance effectiveness, not just governance existence. Specifically, boards should: • Approve and periodically review the AI governance framework. • Provide explicit sign-off that controls are appropriate for the FI's AI activities and risk profile. • Ensure governance considers not only customer harm but also orderly markets and systemic stability impacts where relevant. • Be responsible for ensuring the cross-functional committee (where established) includes sufficient expertise to challenge AI risk decisions and has proportional but independent lines of defence for AI usage **Question 4:** We agree that a dedicated cross-functional committee is essential where AI exposure is material because boards alone generally may not currently have the specialist depth needed. Materiality assessment at organisational level should consider: • Criticality of functions affected (including market-facing functions). • Market impact (disorderly markets/feedback loops), not only direct customer outcomes. • Reliance on AI outputs in decision chains (human-in-the-loop vs human-out-of-the-loop, and degree of automation). • System/model complexity and interaction risk across interconnected workflows and venues. With reference to clause 4.25 in the guidelines, we note a concern that MAS's definitions of "material change" are incorrect for some AI systems; specifically, the implication that model retraining is always a 'minor change' and hence may not be material (see our answer to Q9 below). We also note the possibility that the retraining might be done by a third party which implies a need for such third parties to notify their clients of what they're doing and, potentially, give them the ability to test before deployment. **Question 5:** We support this proposal, noting that it would be helpful to have a taxonomy of AI usage and clear definition of terms (noting clause 4 in the annex provides some useful pointers and also our response to Q2) to help with the identification process and introduce an element of standardization. We also recommend MAS make data governance and provenance the "first-order" control; identification is not simply "is it AI," but "what data, what model, what system, what use case, and what dependencies." Key additions for consideration: • Clear accountability for data management (including where responsibilities sit across IT and outside of IT and across the business). • Data provenance, quality, leakage controls, and explicit treatment of confidential/client data constraints.

		• Suggestion of using disclosure pattern (internal and regulator-facing) aligned to "food label" style transparency, adapted for finance: what data was used, training approach, evaluation results, mitigations, limitations, intended use, and user-data protection. This approach is consistent with the disclosure concept described in the IMDA/AI – https://aiverifyfoundation.sg/wp-content/uploads/2024/06/Model-AI-Governance-Framework-for-Generative-AI-19-June-2024.pdf - p13 • Data Provenance – such as C2PA as noted in our response to Q2. **Question 6:** We support this proposal and refer to our responses to Q2 and Q5 on the need for taxonomies to assist with consistent categorisation. We note that the implied inventory granularity is significantly beyond typical algo inventories such as EU's RTS 6. To be usable, the inventory must capture: • Identifies and documents connections between models, systems and use cases with use of clear and standardised taxonomies where appropriate. • Decision-support tools (research copilots, optimisation, surveillance/risk tooling), not only execution algorithms. • Third-party dependencies and ICT components (consistent with DORA-style third-party risk thinking. • A lifecycle view (pre-build through deployment, monitoring, and retirement). We note that work is underway to extend the existing published FIX Algo Certification Lifecycle for AI-specific needs, such as pre-model data controls, continuous monitoring, and drift management. For the inventory to be effective, procedures and controls need to be put in place to ensure that is it not possible to use AI systems that not on the inventory, and that the inventory is updated in a timely manner when AI systems are modified, replaced or retired. **Question 7:** We support Impact, Complexity and Reliance as a minimum set but recommend adding at least two further dimensions that are critical in practice for AI in market contexts: • Interconnectedness/contagion potential: ability for local failures (or exploitation) to propagate rapidly across venues, participants, and dependent systems. (This reflects the transcript emphasis on nonlinear feedback loops and instability risk.) • Change sensitivity/nonlinearity: small parameter/data changes can yield materially different behaviour, especially in ML systems, which must influence materiality thresholds and retesting requirements. **Question 8:** We agree with MAS's lifecycle framing, but recommend strengthening the minimum expectations around testing, change management, and ongoing monitoring, for example similar to EU RTS 6-style disciplines: • Testing expectations are extremely broad; many firms lack mature frameworks for AI stress testing and model evaluation. • Pre-deployment testing is necessary but insufficient for what is required: controls must include continuous monitoring for drift, degradation, and emergent behaviour. • Change management must treat retraining/fine-tuning and data updates as potentially material and trigger risk-based revalidation.

		• Transparency and inspectability should be emphasised: preserve rich internal audit trails (e.g., full messaging/decision context and appropriate internal tagging) so regulators can inspect behaviour post-incident. • Lifecycle governance should extend RTS 6 concepts into AI: pre-deployment testing, continuous monitoring, clear criteria for material changes, and re-approval workflows. • Recent regulator action reports have stressed the need for end-to-end testing of algo systems in pre-deployment. • Back testing and simulation environments need to be built/adapted to evaluate AI systems when used for trading. • Contribution to market disorder testing is a step beyond conventional module testing, back testing and conformance testing. Regulators have recommended this be done in Dynamic Testing Systems where multiple algorithms can be tested together with stresses and antagonists for realistic interactions. We note that the FIX Trading Community has undertaken significant work on the topics of algorithm testing and certification, some of which may be relevant and helpful to the broader AI topic. We would be happy to provide more information. **Question 9:** We would like to comment on "significant and material change" in clause 4.25. A key gap to address explicitly is the proposed treatment of "minor" vs "material" change. We consider the language implying that retraining may be "minor" is incorrect, particularly for non-linear ML systems. Retraining of such systems can produce logically different models and materially different outcomes. Our recommendation is to revise the definition of "significant/material change" so that retraining, data refreshes, fine-tuning, prompt/policy-layer changes, and infrastructure/latency changes are each evaluated under a clear, risk-based materiality test that can require re-approval and re-testing. We would also like to draw your attention to ISO/IEC 42001 for establishing, implementing, and improving Artificial Intelligence Management Systems with the purpose of ensuring responsible and ethical AI practices. As a global industry association and standards body, the FIX Trading Community is able to draw on the experience of individuals across its membership who are highly experienced in trading and technology governance and risk management. For example, we have, for the last two years, been developing process and messaging standards for the testing and certification (proof of suitability) of trading algorithms and, as mentioned in our response to Q6, believe there are parallels with AI usage. We therefore intend to pursue this work further in the context of AI, and with considerations of the points made in your AI risk management guidelines and would be very happy to discuss this work with you at your convenience. **Question 10:** We believe a fixed 12-month transition window is workable for baseline governance, but risks being obsolete operationally if treated as a one-time compliance project. MAS should position the transition as standing up continuous oversight capability, not "get compliant once." Recommendations:

		• Adopt a phased approach: baseline policies and inventory early; lifecycle controls and advanced testing thereafter, calibrated by materiality. • Encourage iterative supervisory engagement (sandbox-style or structured supervisory check-ins, e.g., as per the UK FCA for AI, and indeed MAS's own regulatory sandbox for digital assets) so requirements can evolve with the technology and observed incidents. • Emphasise rapid detection and containment: focus on preventing "fire spread" (contagion) through transparency, monitoring, and incident handling, rather than assuming prevention is always possible. We stress that innovation in this area is extremely rapid, with technologies emerging in months, not years, and Fis and regulators alike need to be able to adapt regulation and implementation accordingly.
35	Flagright	**Question 1:** We support this risk-based approach. To enhance consistent application, we suggest articulating clearer, tiered risk profiles aligned with factors listed in the Annex. In particular, a structured risk-tiering framework would reduce interpretive variance across institutions. This would strengthen reliability and accountability by promoting consistent interpretation of materiality across institutions, and support consumer protection by reducing the risk of under-controlled high-risk AI use cases being deployed due to inconsistent internal classification. **Question 2:** **Question 3:** **Question 4:** **Question 5:** **Question 6:** We support this requirement and request minimum required attributes be specified to align baseline inventory practices with supervisory expectations. MAS may also wish to clarify triggers for inventory updates, such as changes in model behaviour, data sources, or autonomy level. **Question 7:** We support these dimensions and propose adding data provenance, data sensitivity, and model explainability as minimum dimensions. Data provenance and sensitivity: Data provenance and data sensitivity, particularly the use of personally identifiable information (PII), should be treated as a distinct and minimum dimension in AI risk materiality assessments rather than being subsumed within the broader notion of "impact." As currently framed, the impact dimension encompasses a wide range of considerations, including AI failure, malfunction, and downstream operational effects, which can obscure materially different sources of risk. Functional impact and data-related risk do not always align. An AI system used internally to draft emails or summarise documents may present low impact if it fails, given the

		absence of direct customer or regulatory consequences. However, where such a system consumes, stores, or is trained on PII, the residual risk is materially higher due to privacy, confidentiality, and data protection considerations — regardless of its limited functional impact. In addition, models trained exclusively on an FI's own data generally carry lower privacy, bias, and leakage risk than those trained on external or cross-tenant datasets where data lineage may be less transparent. Distinguishing data provenance and data sensitivity as a separate materiality dimension would enable clearer differentiation between risks arising from AI performance and those arising from the nature and handling of data. Model explainability: We further suggest that explainability be explicitly incorporated as a minimum determinant of materiality, rather than as a control or mitigating factor. The extent to which an AI system can provide traceable, reviewable rationales for its outputs directly affects an FI's ability to exercise effective oversight, validate outcomes, and respond to supervisory or internal challenge. Models that cannot provide meaningful explanations for their outputs, or that update parameters without authorised review, carry elevated residual risk regardless of performance. Such systems warrant higher-risk classification due to limited explainability and the potential to entrench biases or propagate errors from historical data, undermining fairness, accountability, and operational reliability. **Question 8:** We strongly support the lifecycle-based approach in Section 4 and the emphasis on proportionate application of controls. We set below observations on areas where additional clarity would facilitate more consistent implementation. Data management: Paragraph 4.5(b) requires FIs consider Representativeness of Data as a data management control. In practice, models trained primarily on external, pooled or pre-trained datasets may not fully reflect the institution-specific operating conditions, transaction patterns or stress scenarios encountered during deployment. By contrast, models trained or fine-tuned on FI-specific datasets may better support lifecycle controls by enabling more reliable pre-deployment testing and ongoing monitoring for drift. Flagright requests MAS clarify supervisory expectations on how representativeness should be assessed and maintained across the AI lifecycle, particularly where external datasets are used as a baseline but operational conditions are FI-specific. We also note that the data privacy expectations in paragraph 4.5(f) does not address AI models that are developed or deployed across multiple unaffiliated institutions. In such settings, FIs may rely on shared model architectures while remaining accountable for ensuring that customer or employee data is not reused, exposed or inferred across institutional boundaries.

		Flagright requests MAS clarify acceptable privacy-preserving controls where FIs use AI models that are pre-trained by third-party providers and subsequently fine-tuned or adapted using FI-specific customer or employee data. Human oversight Paragraph 4.10 does not explicitly articulate expectations for human involvement in high-risk AI use cases that materially affect consumer outcomes. In the absence of such an expectation, there is a risk that AI systems used in consequential decision-making contexts — such as those affecting access to financial services, customer treatment, or compliance outcomes — may be implemented with insufficient human involvement notwithstanding their assessed risk materiality. We therefore suggest that paragraph 4.10 be strengthened to explicitly state that AI use cases assessed as high risk, particularly those with material impact on consumer outcomes, should incorporate meaningful human oversight and should not be deployed as fully automated decision-making processes in order to preserve fairness and accountability. Third-party AI management: Paragraphs 4.11 and 4.12 appropriately emphasise the need for robust management of third-party AI dependencies. Flagright notes that the practical application of these expectations may be challenging for smaller FIs with limited risk, technology and compliance resources. Many third-party AI components operate as opaque systems with limited transparency into model architectures, training data composition, or internal control environments. While larger FIs may have the capacity to commission bespoke reviews or specialist assessments to address these limitations, smaller FIs are often constrained in their ability to independently validate vendor representations or interpret complex assurance reports. We therefore request that the Guidelines explicitly recognise these proportionality constraints. In addition, MAS may wish to provide illustrative examples of acceptable, proportionate approaches such as baseline documentation, assurance artefacts or compensating controls that smaller FIs may reasonably rely upon where deep independent testing is not feasible, supporting proportional resilience, accountability, and broader financial stability. Evaluation and testing: Flagright notes that behavioural testing frequently extends beyond conventional performance metrics and requests that MAS consider providing illustrative examples of evaluation metrics and testing approaches that may be regarded as good practice. Pre-deployment reviews: Paragraph 4.18 requires that, prior to deployment, AI use cases, systems or models be subject to reviews by parties not involved in their development. Flagright notes that the Guidelines do not explicitly state whether this expectation is intended to include, or exclude, third-party assessments that are commissioned or coordinated by AI vendors.

		Many small FIs rely on vendor-commissioned assurance artefacts due to limited internal capacity. However, where such assessments are scoped, commissioned or managed by the vendor, it is unclear whether they would be regarded as sufficiently independent for the purposes of paragraph 4.18 and 4.19, notwithstanding that the assessors themselves were not involved in model development or deployment. Flagright requests that MAS clarify how "independent" should be interpreted in this context, including the extent to which vendor-commissioned third-party assessments may be relied upon as part of pre-deployment reviews. **Question 9:** AI testing: Paragraph 1.4 of the draft Guidelines acknowledges the rapid evolution of AI systems. Against this backdrop, we note that controlled experimentation is essential for FIs to evaluate emerging AI techniques before determining whether they are suitable for production environments. Existing initiatives such as PathFin and the FinTech Regulatory Sandbox could potentially be extended or complemented to support time-bound, limited-scope testing of AI systems. Agentic systems: We observe that the degree to which an AI model can be configured or fine-tuned by the FI affects governance effectiveness. Models that operate as fixed or non-configurable systems may carry elevated residual risk, as limitations on configurability constrain the FI's ability to remediate issues or align model behaviour with approved risk appetite. In relation to agentic AI systems referenced in paragraph 1.2, some models may autonomously retrieve, combine, or synthesise external information. This presents distinct risk considerations — such as verifying the provenance of retrieved information, assessing reproducibility of outcomes, and documenting reasoning pathways — that are not explicitly addressed in the current text. Highlighting these factors within the Guidelines would support more consistent assessment of risks unique to emerging agentic behaviours, strengthening accountability and systemic resilience. Cross-border data flows: With reference to paragraphs 4.5 and 4.11, many AI services rely on cross-border data flows for training, inference, or model updates. We request clarity on assessing residual risks in scenarios such as: - Where PII is sent outside of Singapore to train or fine-tune third-party AI models; and - When the provenance of training data cannot be fully verified, particularly for models trained on aggregated multi-institution or public internet datasets. AI overlay architectures: In relation to AI overlay architectures referenced under the broader definition of AI in paragraph 1.2, governance considerations may differ materially where new AI components are layered on top of existing business systems to augment their outputs. When such AI components operate using inputs or logic that are not aligned with the underlying system, maintaining clear accountability can become challenging. This challenge is amplified where

		the AI components are inherently opaque or not fully interpretable to the FI, such as cases where they autonomously gather, infer, or interpret external information. In these scenarios, evidencing how inputs were validated, how intermediate reasoning was performed, and how outcomes can be reproduced may not be straightforward. We request additional clarity on supervisory expectations for governance of AI overlay architectures. PII: Building on the materiality and data governance considerations discussed in question 8, we further observe that where AI models trained on PII are used in decision-making, traceability of data usage becomes critical to effective governance and accountability. In such cases, we suggest that FIs should be able to evidence both (i) the PII inputs used to generate a specific decision or outcome, and (ii) the categories and sources of PII used to train or fine-tune the underlying model. Without this level of traceability, it may be difficult for FIs to assess fairness, respond to customer challenges, support supervisory review, or demonstrate compliance with data protection and accountability expectations. **Question 10:** Flagright requests clarity regarding the treatment of new AI deployments initiated during the transition window, as well as whether foundational components — such as governance structures under Section 2, materiality assessments under paragraph 3.10, and inventories under paragraphs 3.4–3.7 — are expected to be prioritised earlier within the period.
36	FWD Singapore Pte Ltd	**Question 1:** Could MAS provide guidance on indicative policy elements or themes that FIs often adopt as baseline good practices? **Question 2:** **Question 3:** **Question 4:** **Question 5:** **Question 6:** **Question 7:** 1. It would be helpful if MAS can provide guidance on how FIs commonly calibrate AI controls relative to materiality ratings. This would help promote proportional implementation of the principles. 2. Since human skill degradation is an emerging topic, it would be useful if MAS provide examples of indicators or considerations that FIs could use when assessing this risk. **Question 8:**

		It would be helpful if MAS can provide examples of evidence or rationale that FIs commonly rely on to demonstrate data suitability or representativeness would help support consistent interpretation. **Question 9:** **Question 10:**
37	The Global Centre for Risk and Innovation, The Global Risks Forum, The Global Risks Alliance (joint submission)	**Question 1:** Question 1: 1. Overall position. 1.1 We support MAS' proposal to apply the Guidelines to all FIs. A baseline set of expectations across the sector is necessary to address systemic and correlated AI risks, prevent regulatory arbitrage, and reduce blind spots arising from embedded AI in vendor products and staff use of GenAI tools. 1.2 Proportionality should calibrate the depth and formality of governance, controls, validation, and evidence requirements based on (i) an FI's overall AI exposure at the organisational level, and (ii) the materiality of each AI use case, system, or model. Proportionality should not remove core obligations. 2. Two-level proportionality framework (recommended). 2.1 Organisational-level calibration ("overall AI exposure"). MAS' guidance is strengthened where organisational expectations scale with the extent to which AI is embedded in the FI's material business processes, critical systems, and control functions. 2.2 Use-case-level calibration ("materiality tiering"). Each AI use case/system/model should be assessed and tiered according to its risk materiality; higher tiers attract higher expectations for controls, independent validation, monitoring, and evidence readiness. 2.3 This two-level approach avoids under-calibration (e.g., a small AI estate containing one high-impact system) and improves consistency and supervisory testability. 3. Non-waivable baseline applicable to all FIs (minimum floor). 3.1 MAS should make explicit in paragraph 1.5 and/or the Annex that all FIs, regardless of size or AI maturity, are expected to maintain a minimum set of baseline measures, including: (a) Board- and senior management-approved AI governance policies covering permitted/prohibited AI uses, accountability, and escalation; (b) Defined roles and ownership for AI risk governance and operational management; (c) A consistent process to identify AI use across the organisation, including embedded/vendor AI and staff-facing AI tools; (d) An accurate, maintained inventory of AI use (at least at a minimum field set), with linkages to technology, data, vendor, change, and incident processes; (e) Change management discipline for material changes to AI systems and dependencies, including vendor changes where reliance is material; (f) Incident response readiness for AI-related events, including escalation pathways and containment expectations; and (g) Minimum third-party AI governance measures proportionate to reliance and criticality.

		3.2 These measures are fundamental for MAS' supervisory objectives and ensure that proportionality does not become a basis for excluding essential governance and operational controls. 4. What should scale with proportionality (calibrated expectations). 4.1 MAS may clarify that proportionality can reasonably be applied to: (a) The depth, frequency, and independence of validation and testing; (b) The sophistication and automation of monitoring (e.g., drift detection granularity and response playbooks); (c) The breadth and detail of documentation (e.g., explainability artefacts proportional to reliance and impact); (d) The scope and frequency of adversarial testing and red teaming; (e) The frequency and granularity of management and committee reporting; and (f) The formality of assurance processes, provided minimum evidence expectations are met. 4.2 Proportionality should not dilute expectations for systems that are high impact, highly relied upon, or difficult to control or reverse. 5. Improve supervisory consistency by expressing proportionality as "tiered controls + tiered evidence". 5.1 To reduce interpretive variance and improve examiner operability, MAS may consider expressing proportionality in the Annex as: (a) Tiered control floors (minimum controls by tier), and (b) Tiered evidence readiness (minimum records and artefacts that the FI must be able to produce on request). 5.2 This approach enables MAS to test outcomes without prescribing a single implementation model, and supports consistent benchmarking across FIs. 6. Embedded AI, vendor AI, and "shadow AI" (explicit treatment recommended). 6.1 MAS may clarify within the proportionality guidance that identification, inventory, and governance extend to: (a) AI embedded within third-party systems used by the FI where reliance is material; and (b) staff-facing GenAI tools and AI features that may be enabled by default in enterprise software. 6.2 Where model internals are not accessible, proportionality should operate on a reliance-based basis: contractual safeguards, performance monitoring, incident obligations, change notification discipline, and exit/portability planning. 7. Practical implementation for smaller FIs without lowering standards. 7.1 MAS may explicitly permit smaller FIs to adopt simplified structures (e.g., combined committees, external managed validation/testing, shared templates), provided they can demonstrate the baseline obligations and produce the minimum evidence required for their AI risk profile. 8. Summary. 8.1 Apply the Guidelines to all FIs with a clear non-waivable baseline.

		8.2 Implement proportionality through (i) organisational-level overall AI exposure calibration and (ii) use-case materiality tiering. 8.3 In the Annex, express proportionality as tiered minimum controls and tiered evidence readiness to enhance supervisory testability and consistency. 8.4 Explicitly include embedded/vendor AI and staff-facing AI usage within identification and governance expectations to reduce blind spots and prevent boundary-gaming. **Question 2:** Question 2: 1. Overall position. 1.1 GCRI supports MAS' proposed approach to scope, covering AI use cases, systems and models that are developed, deployed, operated, procured, or relied upon by FIs across business lines and control functions, including AI embedded within third-party products and services. 1.2 To ensure the Guidelines remain robust for current deployment realities—particularly GenAI, retrieval-augmented generation (RAG) and agentic systems—scope should be anchored not only to the model in isolation but to the end-to-end mechanism by which AI affects decisions, actions and outcomes. 2. Recommended scope anchor: Decision-Influence Chain (DIC). 2.1 MAS may wish to add a concise interpretive concept of a Decision-Influence Chain (DIC) that sits alongside "model / system / use case" and provides a practical, testable scope unit. 2.2 Decision-Influence Chain (DIC) refers to the end-to-end set of components, dependencies, configurations, permissions and human handoffs through which an AI model or AI system is designed, deployed, operated, monitored and changed, and by which its outputs influence decisions, actions or outcomes. 2.3 This aligns with MAS' outcomes-based intent and mitigates boundary-gaming by ensuring that key risk-bearing elements—often outside the model weights—are within governance and control scope. 3. Scope coverage (applied where relevant to the AI use case/system/model). 3.1 In practice, the risk and control perimeter for an AI use case typically spans a chain of components. MAS may clarify that, where applicable, scope encompasses: (a) AI models (including foundation models, fine-tuned variants, classifiers, rankers, ensembles and reinforcement learning components); (b) AI systems (applications, orchestration layers, routing logic, guardrails, output filters, policy engines and decision services); (c) AI use cases (the business context and the specific process steps materially influenced by AI outputs); (d) Data pipelines and transformations used for training, tuning, evaluation and ongoing monitoring; (e) Retrieval/grounding components (RAG corpora, vector stores, search indices, knowledge sources and content governance controls); (f) Prompting and context management (system instructions, templates, context windows, memory/state where deployed);

		(g) Tool/API integrations and permissions (actionability, privilege boundaries, containment controls, and any capability to initiate or modify transactions, communications, or records); (h) Human oversight and handoffs (maker-checker controls, approvals, overrides, escalation routes); (i) Post-processing and downstream logic (thresholds, confidence scoring, ranking, business rules, and decision thresholds); and (j) Monitoring, logging, change management and update pathways, including third-party updates that can materially alter behaviour or performance. 4. Third-party and embedded AI (scope should be reliance-based). 4.1 AI embedded in vendor products or outsourced services should fall within scope where the FI relies on it for material decisions, customer outcomes, market integrity, compliance, or operational resilience, even where model internals are not accessible. 4.2 In such cases, supervisory expectations should be framed on a reliance-based basis: suitable contractual safeguards, change notification and update governance, monitoring for performance and anomalies, incident obligations, and exit/portability planning. 5. Staff-facing AI and unmanaged adoption ("shadow AI"). 5.1 For completeness of identification and governance, the scope of AI usage should include staff use of external GenAI tools, unsanctioned copilots, and AI features enabled by default in enterprise software. These can create material confidentiality, conduct, and operational risk despite not being formally "deployed" systems. 6. AI within control functions (explicit inclusion). 6.1 AI used in risk management, fraud/AML, surveillance, compliance testing, cybersecurity detection, and model monitoring should be treated as within scope. Failures in "control AI" can suppress alerts, distort detection, or delay escalation, creating second-order operational and market integrity risks. 7. Scope boundary (clarity and exclusion). 7.1 A clear exclusion for purely deterministic tools—whose outputs are solely based on predefined programming logic or rules and do not learn or infer—is appropriate, provided such tools do not embed AI components. 7.2 Where a system combines deterministic rules with AI components, the AI component remains within scope and the combined operational setup should be governed as a single end-to-end chain for the relevant use case. 8. Summary. 8.1 The proposed scope is supported and should explicitly cover AI used or relied upon by FIs, including embedded and third-party AI where reliance is material. 8.2 A concise Decision-Influence Chain concept would materially improve operability and supervisory testability for GenAI and agentic systems by ensuring that orchestration, retrieval/grounding, tool permissions, and update pathways are captured within the effective scope. 8.3 Inclusion of staff-facing AI usage and control-function AI reduces blind spots and aligns the scope with real-world risk pathways.

		Question 3: Question 3: 1. Overall position. 1.1 The proposed allocation of responsibilities to the Board and senior management is appropriate and necessary. AI risk is enterprise risk: it affects conduct, operational resilience, cyber risk, model risk, third-party risk, market integrity and customer outcomes. 1.2 Governance should be outcomes-based but examiner-operable: MAS should be able to test whether accountability is clear, whether controls are functioning, and whether the FI can evidence decisions and corrective actions. 2. Board responsibilities (minimum expectations). 2.1 Set risk appetite and boundaries. Approve an AI risk appetite statement that includes: (a) where AI may be used; (b) prohibited use categories; (c) tolerance for automation/autonomy; (d) customer-impact and market-integrity guardrails; and (e) minimum human oversight for higher-risk uses. 2.2 Approve the governance framework and tiering method. Approve: (a) the FI's AI governance framework; (b) a materiality/tiering methodology applied consistently across use cases; and (c) the interpretive scope that governance attaches to the end-to-end "decision-influence chain" for relevant AI deployments (i.e., not the model alone). 2.3 Ensure independence and challenge. Require independent review proportionate to materiality (e.g., internal audit coverage, second-line challenge, and—where appropriate—independent validation/testing for high-impact AI). 2.4 Approve risk acceptance and exceptions discipline. Require that material exceptions/waivers are time-bound, recorded, include compensating controls, and are reported with expiry and remediation status. 2.5 Oversee third-party reliance and concentration risk. Require visibility into AI vendor dependence, shared-model concentration, correlated update risk, sub-outsourcing, and exit/portability readiness for material AI dependencies. 2.6 Oversee incident readiness and customer protection. Require clear incident escalation pathways (including Board notification thresholds for severe AI incidents), containment expectations, and customer impact remediation mechanisms. 2.7 Ensure capability and resources. Ensure adequate resourcing, skills, and authority for first line, second line, technology, and internal audit to discharge AI responsibilities, including secure engineering and monitoring capability. 2.8 Require evidence readiness. Require that the FI can produce decision records and supporting evidence for material AI use cases within supervisory request windows, including inventories, change records, validation artefacts, and incident logs. 3. Senior management responsibilities (minimum expectations). 3.1 Implement Board-approved framework. Translate Board risk appetite into operational policies, standards, controls, and accountability across all relevant functions (business, risk, compliance, technology, procurement, data, operations). 3.2 Maintain identification and inventory discipline. Establish robust mechanisms to identify AI (including embedded/vendor AI and staff-facing AI tools) and maintain a complete, accurate inventory linked to change, incident, vendor, data and access-control processes.

		3.3 Operationalise lifecycle controls. Ensure controls apply across the AI lifecycle—design, data, development, testing/validation, deployment, monitoring, change management, decommissioning—scaled by tier. 3.4 Establish monitoring and thresholds. Implement KRIs/KCIs, drift/performance thresholds, and escalation triggers proportionate to tier and criticality; ensure monitoring covers the full decision-influence chain (e.g., retrieval corpora, tool permissions, orchestration/routing, and vendor updates). 3.5 Manage change and "no silent updates" for material AI. Ensure that material changes—including vendor model/API upgrades, retrieval/grounding corpus changes, routing/guardrail changes, and permission changes—are governed, logged, tested as needed, and approved according to tier. 3.6 Third-party governance. Ensure contractual and operational controls for third-party AI, including: disclosure of sub-outsourcing, change notification, incident obligations, audit/assurance rights, data use limits, and exit/portability plans for material dependencies. 3.7 Incident management and drills. Maintain AI incident runbooks, perform containment and recovery drills for high-risk AI, and ensure timely remediation and customer outcomes management. 3.8 Culture and "shadow AI" controls. Establish enforceable staff rules, training, and technical controls to prevent unmanaged AI adoption and to protect confidential data and customer information. 4. Governance operating model and reporting (to reduce variance). 4.1 Designated accountable executive. MAS may explicitly expect a named senior executive accountable for enterprise AI risk management, with authority to enforce inventory, gating, and stop-the-line controls. 4.2 Cross-functional committee trigger. Where overall AI exposure is material (or any high-risk tier exists), senior management should operate a cross-functional AI risk committee with clear terms of reference, including portfolio concentration oversight and exception management. 4.3 Board reporting minima. Board (or delegated Board committee) reporting should include, at minimum: (a) inventory completeness and coverage; (b) distribution of AI use cases by tier and business line; (c) material changes and vendor updates affecting higher tiers; (d) validation/testing outcomes and unresolved issues; (e) incidents/near-misses and corrective actions; (f) exceptions/waivers and expiry status; (g) concentration/coupling and correlated update risk; and (h) capability metrics (skills, resourcing, audit coverage). 5. Clarifications that would strengthen MAS' drafting. 5.1 Accountability follows reliance, not model ownership. For third-party or embedded AI, governance accountability should attach to reliance and criticality, even where internals are inaccessible. 5.2 "Decision-influence chain" as scope for oversight. Board and senior management responsibilities should be understood to cover the end-to-end mechanism by which AI

		influences outcomes (including orchestration, retrieval/grounding, permissions, and updates), as these are frequent sources of failure in GenAI and agentic systems. 5.3 Evidence and correction discipline. Oversight should include a requirement to maintain decision records, change records, and correction logs for material AI so that outcomes are reconstructable and improvements are demonstrable. 6. Summary. 6.1 The Board should set AI risk appetite, approve tiering and governance, ensure independent challenge, oversee concentration and third-party reliance, and require evidence readiness and incident preparedness. 6.2 Senior management should operationalise lifecycle controls, maintain identification and inventory completeness, govern material changes (including vendor updates), ensure robust monitoring and incident response, and prevent unmanaged AI adoption. 6.3 Explicit Board reporting minima and a practical scope anchor to the end-to-end decision-influence chain would materially improve consistency, supervisory testability, and real-world effectiveness of the Guidelines. **Question 4:** Question 4: 1. Overall position. 1.1 The proposal to require a dedicated cross-functional committee where an FI's overall AI risk exposure is material is appropriate and should strengthen accountability, reduce control fragmentation, and improve consistency of AI risk decisions across business lines and functions. 1.2 To maximise supervisory operability and reduce interpretive variance, MAS may consider (i) specifying a clear trigger framework for when exposure is "material", and (ii) setting minimum mandate and reporting expectations for such a committee. 2. When a dedicated cross-functional committee should be required (recommended triggers). 2.1 A committee should be required where any of the following conditions are met: (a) The FI operates or relies on any high-impact AI use case (e.g., AI materially influencing customer outcomes, market integrity, credit decisions, fraud/AML outcomes, risk limits, pricing, trading, or regulatory compliance); (b) The FI deploys GenAI or agentic AI with elevated autonomy, broad tool/API access, or the ability to initiate or modify transactions, communications, or records; (c) AI is used in control functions (risk management, surveillance, AML/fraud, compliance testing, cybersecurity detection, model monitoring) in a way that could suppress alerts, distort detection, or delay escalation; (d) The FI exhibits material third-party reliance on external AI providers, including concentration on a small number of shared models/endpoints or high exposure to correlated updates; (e) The FI's AI estate exceeds a specified portfolio threshold (recommended in section 4 below) based on aggregated impact, reliance and complexity. 2.2 MAS may allow proportional implementation for smaller FIs (e.g., a combined committee or adapted structure), provided mandate, accountability, and evidence readiness expectations are met.

		3. Committee purpose and minimum mandate (examiner-operable). 3.1 The committee should function as the enterprise "control plane" for AI risk, ensuring that AI is governed consistently and that material risks are identified, measured, controlled, and corrected. 3.2 Minimum mandate should include: (a) Approving and maintaining the FI's AI taxonomy and identification standards across the enterprise; (b) Owning the AI inventory completeness and accuracy assurance process; (c) Approving tiering/materiality determinations, re-tiering decisions, and risk acceptance for material AI uses; (d) Governing material changes across the end-to-end chain through which AI influences outcomes (including orchestration, retrieval/grounding, permissions, and vendor updates); (e) Overseeing lifecycle control effectiveness (testing, validation, monitoring, drift controls, incident readiness); (f) Managing exceptions/waivers discipline (time-bounded, recorded, with compensating controls and expiry tracking); (g) Overseeing third-party AI risk, including contract minima, sub-outsourcing transparency, change notification, and exit/portability plans; (h) Monitoring portfolio concentration/coupling and correlated update risk; and (i) Reviewing incidents, near-misses, and corrective action closure, including containment drill outcomes for higher-risk AI. 4. How "overall AI risk exposure" should be assessed at organisational level (recommended approach). 4.1 Organisational exposure should be assessed using a portfolio-based assessment that aggregates AI risk across all use cases/systems/models and their dependencies. This should be measurable, repeatable, and evidence-supported. 4.2 A practical approach is to assess overall exposure across at least five organisational dimensions: (a) Portfolio impact footprint ○ Volume and criticality of AI use cases affecting customers, financial outcomes, market integrity, or regulatory compliance; ○ Extent of AI influence on material decisions (direct decisioning vs advisory). (b) Portfolio reliance and critical dependency ○ Degree to which business continuity depends on AI outputs; ○ Whether manual fallback exists and whether it is operationally viable under stress. (c) Portfolio complexity and change velocity ○ Breadth of AI architectures and heterogeneity of models/systems;

		○ Frequency of changes, retraining, prompt/guardrail updates, retrieval corpus updates, routing changes, and vendor version updates. (d) Portfolio autonomy and actionability ○ Presence of GenAI/agentic systems with tool/API access; ○ Blast radius potential based on permissions, integration points, and speed of execution. (e) Portfolio concentration/coupling and third-party dependence ○ Dependence on shared vendor models/endpoints; correlated update risk; ○ Sub-outsourcing opacity; exit/portability constraints; data residency and handling constraints. 4.3 MAS may consider requiring an FI to maintain a "portfolio exposure scorecard" (qualitative + quantitative) that is reviewed at least annually and upon material changes, and that supports supervisory discussion. 5. Materiality determination (recommended clarity). 5.1 To reduce debate on "material exposure", MAS may define that exposure is deemed material where: (a) the FI has any AI use case classified in the highest materiality tier; or (b) AI is used in critical control functions or critical customer outcome pathways; or (c) portfolio concentration/coupling indicates that failure of a shared vendor model, endpoint, or control AI could materially impair the FI's operations, risk management, compliance, or customer outcomes. 6. Governance and reporting expectations. 6.1 The committee should have representation from relevant functions at minimum: business owners (first line), risk management (second line), compliance, technology/cybersecurity, data governance, operations, procurement/vendor management, and internal audit (as observer or periodic reviewer). 6.2 It should operate under clear terms of reference, with documented decisions and evidence trails sufficient for audit and supervisory review. 6.3 Board (or delegated Board committee) reporting should include inventory completeness, tier distribution, key model/system risks, material changes (including vendor updates), incidents/near-misses, exceptions, and concentration/coupling risk. 7. Summary. 7.1 The committee requirement is supported and should be triggered by high-impact AI use, material third-party reliance, GenAI/agentic autonomy, control-function AI reliance, or portfolio exposure thresholds. 7.2 Organisational exposure should be assessed using a portfolio method spanning impact, reliance, complexity/change velocity, autonomy/actionability, and concentration/coupling (including correlated updates).

		7.3 Clear mandate, decision logging, and reporting minima will make the requirement operationally effective and supervisory testable while remaining proportionate across FI profiles. **Question 5:** Question 5: 1. Overall position. 1.1 The proposal is supported. Consistent identification of AI across business and functional areas is a foundational control: without reliable identification, FIs cannot apply governance, tiering/materiality assessments, lifecycle controls, monitoring, incident response, or third-party safeguards in a complete and consistent manner. 1.2 Identification should be designed to prevent common failure modes, including (i) embedded AI in vendor products being missed, (ii) staff use of GenAI tools ("shadow AI") bypassing governance, and (iii) model-only scoping that excludes orchestration, retrieval/grounding, tool permissions, and update pathways that often drive real-world risk. 2. Definitions and scope standard (minimum). 2.1 FIs should adopt a single enterprise AI taxonomy aligned to MAS' "model / system / use case" terms and apply it consistently across all lines of business and functions. 2.2 For operational completeness, identification should extend to the end-to-end chain through which AI influences outcomes (e.g., orchestration/routing, retrieval/grounding, prompt/guardrail layers, tool/API integrations and permissions, post-processing logic, monitoring components, and third-party updates) where these are integral to the AI use case. 2.3 FIs should clearly state boundary rules for exclusions (e.g., purely deterministic rule-based tools), while ensuring hybrid rule+AI systems remain in scope for the AI component. 3. Identification criteria (what triggers an "AI" classification). 3.1 The FI should define objective criteria to identify AI, including whether a component: (a) learns from data or adapts; (b) performs probabilistic inference, ranking, classification, prediction, generation, or optimisation; (c) produces outputs that influence decisions, actions, customer outcomes, market outcomes, or controls; (d) uses GenAI capabilities (text/image/code generation, summarisation, translation, reasoning); (e) includes agentic characteristics (task planning, tool use, multi-step action execution); or (f) is embedded as an AI feature in third-party software, platforms, APIs, or managed services relied upon by the FI. 4. Processes to ensure consistent enterprise-wide identification (minimum). 4.1 End-to-end coverage points. Identification should be embedded into the FI's standard operating processes at multiple entry points, including: (a) procurement and vendor onboarding (including SaaS feature reviews and "AI-by-default" features); (b) architecture and design review;

		(c) model/system onboarding and release gating; (d) SDLC/ML lifecycle controls (development, testing, deployment); (e) information security and data governance approvals; (f) change management (including vendor upgrades, prompt/guardrail updates, retrieval corpus updates, routing/permission changes); (g) incident management and operational risk events; and (h) periodic discovery reviews (to detect drift in the AI estate and unmanaged adoption). 4.2 Three lines integration. Identification responsibilities should be explicit across the three lines of defence: (a) first line identifies and registers AI at inception and maintains completeness in the business process; (b) second line defines criteria, challenges completeness, and validates tiering and control assignment; (c) internal audit periodically tests completeness and effectiveness. 5. Robust systems (what "system-supported" identification should mean). 5.1 Identification should not rely on ad-hoc declarations. It should be supported by systems that provide consistency, auditability, and linkage to control operations. 5.2 At minimum, FIs should maintain a system of record (AI register) integrated with: (a) CMDB / application inventory; (b) data catalog and lineage tooling; (c) vendor register and contract repository; (d) IAM and privilege management (particularly for tool/API permissions); (e) change management systems; (f) incident and operational risk systems; and (g) model monitoring and logging infrastructure. 5.3 Where feasible, automated discovery controls may be used to detect AI usage signals (e.g., approved model endpoints, AI API consumption, GenAI tool access patterns, agent tool calls), with escalation and remediation workflows. 6. Coverage of embedded AI and third-party AI (explicit treatment). 6.1 Identification criteria and processes should require procurement and technology functions to flag AI features embedded in vendor products, including upgrades that introduce new AI capabilities. 6.2 Where internals are not available, identification should still register the AI dependency at the system and use-case level, including reliance classification, concentration/coupling indicators, and change notification pathways. 7. Coverage of staff-facing AI and "shadow AI". 7.1 Identification must cover staff use of external GenAI tools and copilots to prevent data leakage, conduct risks, and unmanaged operational dependency. 7.2 FIs should adopt clear internal rules and technical controls: approved tool lists, data handling constraints, access controls, monitoring, and mandatory reporting/escalation for new AI tool adoption. 8. Summary.

		8.1 The proposal is supported and should be implemented through clear definitions, objective identification criteria, and embedded identification processes across procurement, architecture, SDLC, change, incident and periodic discovery reviews. 8.2 "Robust systems" should mean a system-of-record AI register integrated to CMDB, data lineage, vendor management, IAM, change and incident tooling, enabling auditability and consistent downstream control application. 8.3 Explicit coverage of embedded vendor AI and staff-facing "shadow AI" is essential to achieve the policy intent and reduce supervisory blind spots. **Question 6:** Question 6: 1. Overall position. 1.1 The proposal is strongly supported. An accurate, up-to-date inventory is the foundational "system of record" for AI risk governance and is necessary for consistent tiering/materiality assessment, lifecycle controls, monitoring, change management, incident response, third-party governance, and supervisory testability. 1.2 In practice, inventory obligations should extend beyond "models" to the deployed AI use cases and systems, including key dependencies that materially affect behaviour (e.g., orchestration/routing, retrieval/grounding corpora, tool/API permissions, and vendor update pathways) where relevant. 2. Inventory as a governed system of record (minimum expectation). 2.1 The inventory should be a formally governed register with defined ownership, data quality standards, update SLAs, and periodic completeness attestations. 2.2 The inventory should be linked to the FI's broader control environment (technology, data, vendor, security, change, incident), such that inventory entries drive downstream control obligations and monitoring requirements. 2.3 The inventory should support evidence readiness—i.e., the FI can produce accurate inventory extracts and supporting records within supervisory request windows. 3. Minimum inventory coverage (what must be captured). 3.1 At minimum, the inventory should cover all AI usage across: (a) business decisioning (e.g., credit, underwriting support, pricing support, marketing and customer engagement); (b) trading, treasury, and financial risk functions (where applicable); (c) operational processes (service, collections, onboarding, dispute handling); (d) compliance and surveillance (AML/fraud, conduct, monitoring, regulatory reporting support); (e) cybersecurity and technology operations (detection, response, automation); (f) internal support functions (HR, finance operations) where AI influences decisions or processing; and (g) third-party and embedded AI dependencies relied upon by the FI. 4. Minimum inventory data fields (examiner-operable). 4.1 MAS may consider specifying a minimum field set to reduce variance. A practical minimum includes: Identity and ownership

		(a) unique inventory identifier; (b) business owner; (c) technical owner; (d) accountable senior executive; (e) vendor owner (if applicable). Purpose and scope (f) AI use case description; (g) affected business process and decision points; (h) customer/market impact touchpoints; (i) degree of automation/actionability (advisory vs decisioning vs agentic execution). Model/system descriptors (j) model type and family (where applicable); (k) model/system name; (l) version; (m) hosting/deployment model (cloud/on-prem/vendor API); (n) key dependencies (orchestration/routing components, guardrails/post-processing). Data, grounding and permissions (o) key data sources and sensitivity classification; (p) retrieval/grounding sources (e.g., RAG corpus identifiers, knowledge source governance) where applicable; (q) tool/API integrations; (r) permission scope and containment controls. Materiality, controls and assurance (s) tier/materiality rating and rationale; (t) required controls baseline per tier; (u) validation/testing status and dates; (v) monitoring approach and key thresholds (drift/performance); (w) fallback/override arrangements and operational viability. Third-party and concentration (x) vendor/service provider and sub-outsourcing visibility; (y) reliance classification; (z) concentration/coupling indicators (shared model endpoints, correlated update dependencies); (aa) exit/portability status. Change and incident history (ab) change log references (including vendor update events); (ac) incident/near-miss references and corrective actions; (ad) exceptions/waivers and expiry dates; (ae) decommissioning status where relevant. 5. Update discipline and “no silent updates” for material AI. 5.1 FIs should maintain defined update SLAs and triggers to ensure timeliness, including mandatory updates when any of the following occur: (a) new AI use case/system/model onboarding; (b) material change to model/system version, architecture, or deployment; (c) prompt/guardrail changes for GenAI; (d) retrieval/grounding corpus changes; (e) orchestration/routing logic changes; (f) tool/API permission changes; (g) vendor model/API upgrades or capability expansions; (h) monitoring thresholds or validation baselines change; or (i) significant incidents, near misses, or supervisory findings.

		5.2 For higher materiality tiers, the FI should ensure "no silent updates" in practice: vendor updates and dependency changes should be logged, assessed for impact, and governed proportionately. 6. Integration with enterprise systems (robustness and auditability). 6.1 The inventory should be integrated with, or reconciled against, relevant enterprise systems to reduce reliance on manual maintenance, including: (a) CMDB/application portfolio; (b) data catalog/lineage tooling; (c) vendor/contract repository; (d) IAM/privileged access management; (e) change management; (f) incident and operational risk systems; and (g) monitoring/logging platforms. 6.2 Where feasible, automated detection (e.g., AI endpoint/API usage, GenAI tool access telemetry, agent tool calls) can provide completeness checks and "shadow AI" discovery signals. 7. Governance, assurance, and supervisory readiness. 7.1 MAS may consider expecting periodic management attestations on inventory completeness and accuracy, with independent testing by internal audit on a risk-based cycle. 7.2 The inventory should support supervisory engagement by enabling prompt retrieval of: tier distribution, high-impact use cases, third-party dependencies, material changes, incidents, exceptions, and corrective action status. 8. Summary. 8.1 The inventory requirement is supported and should be treated as a governed system of record integrated with the FI's control environment. 8.2 A minimum field set, defined update triggers, and explicit treatment of vendor updates and key dependencies will materially improve consistency, risk coverage, and supervisory testability. 8.3 Integration with enterprise systems and telemetry-based completeness checks will reduce manual errors and improve ongoing accuracy, including for embedded and "shadow AI" usage. **Question 7:** Question 7: 1. Overall position. 1.1 The proposed risk dimensions of impact, complexity and reliance are appropriate as the minimum core triad for AI materiality assessments. They align to the principal determinants of harm, controllability, and dependency in FI environments. 1.2 To strengthen real-world effectiveness—particularly for GenAI, RAG and agentic systems—materiality assessments should also capture a small set of additional dimensions that consistently drive incident severity, supervisory concern, and systemic risk. 2. Comments on the proposed core dimensions (retain as minimum).

		2.1 Impact should capture the magnitude and distribution of potential adverse outcomes (customers, market integrity, safety and soundness, compliance breaches, operational resilience), including both direct and second-order harms (e.g., suppressed alerts, delayed escalation, or compounding errors). 2.2 Complexity should capture technical and operational complexity across the end-to-end deployment chain (including data pipelines, orchestration/routing, retrieval/grounding, guardrails, tool integrations, and frequency of changes), as complexity materially increases control fragility and the probability of undetected failure. 2.3 Reliance should capture the degree of dependency on AI outputs for decisioning and control effectiveness, including the availability and viability of manual fallback under stress and the extent of "control substitution" (i.e., AI replacing human or deterministic controls). 3. Additional risk dimensions recommended for inclusion (minimal set, high explanatory power). 3.1 Autonomy and actionability (agent/tool power). 3.2 Whether AI can initiate or modify transactions, communications, customer records, risk limits, or operational actions via tool/API access; whether it can plan and execute multi-step tasks. 3.3 This dimension is decisive in GenAI/agentic deployments because it governs speed, blast radius, and containment difficulty. 4. Reversibility and containment (recoverability). 4.1 The ability to halt, isolate, and rollback effects within defined timeframes, including kill-switch capability, privilege revocation, and safe fallback operations. 4.2 This dimension translates directly into operational resilience and incident severity. 5. Data sensitivity and confidentiality handling. 5.1 Whether AI touches or can infer sensitive information (customer PII, confidential supervisory information, privileged trading/risk data, security telemetry), including leakage risk through prompts, retrieval corpora, outputs, and logs. 5.2 This dimension is central for GenAI and RAG due to inadvertent disclosure pathways. 6. Observability, traceability, and auditability (reconstructability). 6.1Whether the FI can reconstruct decisions and outcomes with sufficient logs and lineage across the full chain (inputs, prompts/context, retrieval sources, model version, routing, tool calls, human approvals, post-processing). 6.2This dimension is essential for supervisory testability, internal accountability, and effective remediation. 7. Third-party opacity and change dependence.

		7.1Degree to which model behaviour is influenced by vendor-controlled updates, limited transparency, sub-outsourcing, and restricted audit rights; includes portability and exit constraints. 7.2 This is distinct from "reliance": it captures controllability and change risk under third-party governance. 8. Concentration, coupling, and correlated update risk (systemic dependency). 8.1 The extent to which multiple critical functions depend on the same model endpoint, vendor, shared service, or common retrieval corpus—and the risk that a single update or outage causes widespread correlated failure. 8.2 This dimension is increasingly material as FIs adopt the same foundation model providers and shared AI tooling. 9. How MAS could keep the framework proportionate and operable. 9.1 MAS can retain the triad (impact, complexity, reliance) as the minimum core and require FIs to incorporate the additional dimensions above either: (a) as explicit sub-factors under the triad (e.g., autonomy under impact/complexity; containment under reliance/impact), or (b) as a small set of supplementary dimensions required for GenAI/agentic and other high-impact tiers. 9.2 This preserves proportionality for low-impact use cases while ensuring the framework is fit-for-purpose for the most consequential deployments. 10. Summary. 10.1 The triad of impact, complexity and reliance should be retained as the minimum. 10.2 MAS should add (or explicitly require consideration of) autonomy/actionability, reversibility/containment, data sensitivity, observability/auditability, third-party opacity/change dependence, and concentration/coupling (correlated update) risk. 10.3 These additions materially improve predictive power for real-world incidents and strengthen supervisory testability, particularly for GenAI, RAG and agentic systems. **Question 8:** Question 8: 1. Overall position. 1.1 The proposed lifecycle-wide approach is supported. Effective AI risk management requires controls that operate continuously—from design and development through deployment, monitoring, change, incident response and decommissioning—rather than point-in-time documentation. 1.2 The "relevance" and "proportionality" concepts are appropriate, but should be made examiner-operable by clarifying: (i) which controls are baseline and non-waivable, (ii) how relevance is determined, and (iii) how higher-risk AI (including GenAI/RAG/agentic systems) triggers additional controls focused on orchestration, permissions, and update pathways that are frequent sources of failure.

		2. Lifecycle control objectives (what controls should achieve). 2.1 Across the lifecycle, controls should ensure: (a) clear accountability and governed use; (b) controlled data and model inputs/outputs; (c) demonstrable safety, robustness and performance prior to deployment; (d) ongoing monitoring for drift, degradation and anomalies; (e) controlled and logged change with impact assessment; (f) incident readiness, containment and timely remediation; and (g) evidence readiness for audit and supervisory review. 3. Non-waivable baseline controls (apply to all FIs and all AI in scope). 3.1 MAS may clarify that certain controls are always required regardless of proportionality, including: (a) identification and inventory as a system of record; (b) ownership and accountability assignment; (c) governance of third-party AI reliance and embedded AI; (d) change management discipline for material changes; (e) incident escalation and response arrangements; (f) cybersecurity baseline controls appropriate to the FI's environment; and (g) minimum recordkeeping sufficient to reconstruct key decisions and changes for material AI. 4. Key lifecycle stages and expected controls (proportionate by materiality tier). 4.1 Strategy, scoping and use-case approval (a) documented purpose, intended users, and decision points influenced; (b) materiality tiering based on impact/complexity/reliance and other relevant dimensions (autonomy, reversibility, data sensitivity, observability, concentration); (c) defined risk acceptance thresholds and prohibited use categories; (d) defined human oversight model (advisory vs decisioning vs agentic execution). 4.2 Data governance and preparation (a) data quality controls and lineage documentation proportionate to tier; (b) privacy, confidentiality and handling controls, including prompt and output handling for GenAI; (c) governance of training, tuning, evaluation, and grounding data (RAG corpora) and their update pathways; (d) bias/fairness assessment proportionate to customer impact and regulated decision contexts. 4.3 Model/system design and development (a) secure development practices; (b) documentation of model/system architecture and dependencies; (c) control design for the full operational chain (orchestration/routing, guardrails, post-processing, retrieval/grounding, tool/API integrations); (d) segregation of duties and appropriate peer review. 4.4 Testing, validation and pre-deployment gating (a) performance evaluation against defined metrics and acceptance criteria;

		(b) robustness testing (including stress testing, adversarial testing where relevant); (c) explainability and transparency artefacts proportionate to reliance and impact; (d) validation of monitoring and alerting effectiveness; (e) independent validation/review for high-impact tiers; (f) documented fallback and manual override procedures, tested for viability. 4.5 Deployment and release management (a) controlled release with versioning and rollback capability; (b) environment controls (access management, secrets management, logging); (c) segregation of production and testing; (d) "no silent updates" expectations for material AI dependencies, including vendor model/API upgrades and retrieval corpus changes. 4.6 Monitoring, performance management and periodic review (a) continuous monitoring for drift, degradation, bias and anomalies; (b) defined thresholds, escalation triggers and remediation playbooks; (c) periodic re-validation and re-tiering triggers; (d) monitoring of third-party service reliability and behavioural changes. 4.7 Change management (including vendor and dependency changes) (a) formal change assessment for material updates, including prompts/guardrails, retrieval/grounding corpora, routing logic, tool permissions and vendor updates; (b) tier-based testing requirements prior to release; (c) recorded approvals and audit trail of changes; (d) expiry-based exception/waiver discipline with compensating controls. 4.8 Incident management and resilience (a) AI incident classification and escalation pathways (including customer harm and market integrity triggers); (b) containment controls (kill switch, privilege revocation, isolation, throttling); (c) post-incident review, root cause analysis, and documented corrective actions; (d) periodic drills for higher-risk tiers (including correlated vendor failure scenarios). 4.9 Decommissioning and retirement (a) controlled decommissioning with data retention/secure disposal rules; (b) replacement risk assessment where AI is substituted; (c) preservation of necessary records for audit and supervisory purposes. 5. Key areas for relevance assessment (what FIs should assess and apply proportionately). 5.1 MAS' "apply what is relevant" approach can be strengthened by specifying that relevance should be assessed at least across: (a) customer/market impact; (b) reliance and ability to operate without AI (fallback viability); (c) complexity and change velocity; (d) autonomy/actionability (tool permissions, agent behaviour); (e) data sensitivity and confidentiality handling; (f) observability/auditability (ability to reconstruct); (g) third-party dependence and opacity; and

		(h) concentration/coupling and correlated update risk. 6. GenAI/RAG/agentic controls (recommended explicit additions). 6.1 MAS may explicitly call out additional controls for GenAI, RAG, and agentic systems where relevant, including: (a) prompt, system instruction and context governance (versioning, approval, handling controls); (b) retrieval/grounding governance (approved corpora, provenance controls, update gating); (c) tool/API permission management (allowlists, least privilege, segregation, high-risk action gating); (d) agent containment (rate limits, budget limits, sandboxing, kill switch, step-level approvals for high-risk actions); (e) output risk controls (hallucination containment, unsafe content filters where relevant, customer communication approvals); and (f) "no silent updates" across the decision-influence chain for higher tiers. 7. Making proportionality testable (tiered controls and tiered evidence). 7.1 MAS may consider framing lifecycle expectations in the Annex as tiered minima: for each tier, specify minimum controls and minimum evidence artefacts (e.g., validation records, monitoring thresholds, change logs, incident playbooks). 7.2 This preserves outcomes-based flexibility while enabling consistent supervisory review. 8. Summary. 8.1 The lifecycle-wide standards, processes and controls are supported; controls should operate continuously across design, build, test, deploy, monitor, change, incident and retirement. 8.2 MAS should clarify non-waivable baseline controls and require relevance/proportionality to be applied through a tiering method that is consistent and evidence-supported. 8.3 Explicit coverage of GenAI/RAG/agentic systems (orchestration, retrieval, tool permissions, containment, and update governance) will materially strengthen the Guidelines' effectiveness and reduce blind spots. 8.4 Framing proportionality as tiered controls plus tiered evidence artefacts in the Annex will improve operability, reduce variance, and increase supervisory testability. **Question 9:** Question 9: 1. Overall position. 1.1 The Guidelines are comprehensive and appropriately outcomes-based. The remaining opportunities are primarily in areas that reduce supervisory variance, address GenAI/agent operational realities, and strengthen systemic resilience in a sector increasingly exposed to shared-model concentration and correlated update risks. 2. Reduce variance through "minimum evidence readiness" expectations. 2.1 MAS may consider adding an explicit expectation that FIs maintain evidence readiness for material AI, i.e., the ability to produce, within supervisory request windows, a minimum

		set of records demonstrating governance, validation, monitoring, change control, and incident response. 2.2 Evidence readiness can be defined without requiring disclosure of sensitive model internals: the focus is on decision records, control operation receipts, and reconstructability of material outcomes. 3. Explicit treatment of GenAI, RAG and agentic systems (control-plane specificity). 3.1 While GenAI is in scope, MAS may strengthen the Guidelines by explicitly referencing common failure pathways specific to GenAI/RAG/agents, including: (a) prompt and system-instruction governance (versioning, approval, handling constraints); (b) retrieval/grounding governance (approved corpora, provenance, update gating); (c) tool/API permission management (allowlists, least privilege, segregation, high-risk action gating); (d) agent containment (rate limits, budget limits, sandboxing, kill switch, step-level approvals for higher-risk actions); and (e) output controls for customer-facing uses (accuracy checks, disclosure, escalation for uncertainty). 3.2 These measures directly address "decision-influence chain" risks that are frequently outside the model weights. 4. Portfolio concentration, coupling and correlated update risk (systemic dependency). 4.1 MAS may consider including an explicit expectation for FIs to monitor concentration/coupling in their AI estate, particularly: (a) reliance on a small set of foundation model providers or shared endpoints; (b) dependence on common orchestration layers or shared retrieval sources; and (c) correlated update risk where a single vendor update can affect multiple critical functions simultaneously. 4.2 This is increasingly relevant as sector-wide adoption of a few providers can create systemic fragility not visible in single-system assessments. 5. Third-party AI contractual minima (greater specificity). 5.1 MAS may strengthen third-party guidance by stating minimum contractual expectations for material AI dependencies, including: (a) transparency on sub-outsourcing and material dependency chains; (b) change notification for model/API/version/capability updates; (c) data use limits, isolation boundaries, and confidentiality handling; (d) incident notification clocks and cooperation obligations; (e) audit/assurance rights proportionate to reliance; and (f) portability/exit planning and discontinuation safeguards. 5.2 This supports reliance-based governance where internals are not accessible. 6. Clear "exceptions/waivers" discipline with expiry and compensating controls. 6.1 MAS may consider explicitly requiring that exceptions/waivers to controls are: (a) time-bounded with expiry; (b) documented with rationale and compensating controls; (c) visible to the cross-functional committee where required; and (d) escalated to the Board (or delegated Board committee) where material AI is involved.

		6.2 This reduces the risk that exceptions become permanent control gaps. 7. Incident management enhancements specific to AI. 7.1 MAS may consider including guidance on AI incident typologies and severity triggers (e.g., customer harm, conduct breach, market integrity concern, control suppression, significant data leakage, agent mis-execution). 7.2 For higher tiers, require periodic containment and recovery drills (including vendor outage/correlated update scenarios), and ensure "stop-the-line" authority is clearly assigned. 8. Handling controls for confidential/sensitive data in GenAI contexts. 8.1 MAS may clarify expectations for handling controls covering: (a) prompt and output retention policies; (b) restrictions on entering sensitive information into external tools; (c) safeguards for retrieval corpora that contain sensitive or regulated data; and (d) segregation of supervisory, risk and security data used in control AI. 8.2 This is a common risk pathway even in otherwise low-impact GenAI usage. 9. Human oversight quality (beyond "human-in-the-loop"). 9.1 Where human oversight is required, MAS may clarify that effective oversight includes: (a) clear decision rights and escalation paths; (b) operator competence and training; (c) manageable workload and alert fatigue controls; and (d) recorded rationale for overrides and approvals for material decisions. 9.2 This prevents "nominal" oversight that is ineffective in practice. 10. Interoperability and standardisation for supervisory efficiency (optional). 10.1 MAS may consider encouraging standardised inventory fields and evidence artefact structures (templates) to reduce reporting burden and improve comparability across FIs, while remaining outcomes-based. 10.2 This can be phased, starting with minimum fields for inventory and minimum evidence sets for high-tier AI. 11. Summary. 11.1 The most valuable additions are: (i) minimum evidence readiness for material AI; (ii) explicit GenAI/RAG/agent control-plane measures; (iii) portfolio concentration/coupling and correlated update monitoring; (iv) more specific third-party contractual minima; (v) time-bounded exceptions discipline; and (vi) AI-incident typologies, containment drills, and clear stop-the-line authority. 11.2 These enhancements improve real-world effectiveness and supervisory testability without undermining proportionality or innovation. **Question 10:** Question 10: Proposed transition period of 12 months 1. Overall position. 1.1 A 12-month transition period is broadly appropriate and feasible, provided MAS adopts a phased implementation expectation that prioritises (i) identification and inventory

		completeness, (ii) materiality tiering, and (iii) control effectiveness for the highest-risk AI first. 1.2 A phased approach reduces execution risk, avoids "paper compliance", and aligns supervisory effort to the most consequential AI deployments, including GenAI/RAG/agentic systems and material third-party AI dependencies. 2. Recommended phasing within the 12-month period (examiner-operable). 2.1 MAS may consider articulating the transition as three implementation waves with clear exit criteria: Wave 1 (Months 0–3): Establish control foundations (a) Adopt enterprise AI definitions and identification criteria aligned to MAS' model/system/use case terms and an end-to-end scope anchor for material AI deployments; (b) Assign accountable owners (business, technology, second line) and designate senior executive accountability; (c) Establish the AI inventory as a governed system of record with minimum fields and update triggers; (d) Commence enterprise discovery to identify embedded/vendor AI and unmanaged staff-facing AI usage; (e) Define materiality/tiering methodology and initial tiering for known AI use cases; and (f) Implement baseline staff rules for GenAI tools (approved tools, prohibited data handling, escalation). Wave 2 (Months 3–6): Operationalise tiering and high-risk controls (a) Complete baseline inventory coverage across business lines and functional areas, linked to CMDB, vendor register, change and incident systems; (b) Establish cross-functional committee arrangements where overall AI exposure is material (or where high-risk tier exists), with defined mandate and decision logging; (c) Implement tier-based lifecycle gates for higher tiers: validation/testing, monitoring thresholds, change management, incident playbooks; (d) Put in place third-party AI contractual minima for material dependencies (change notification, incident clocks, audit/assurance rights proportionate to reliance, sub-outsourcing visibility, data use limits, exit planning); (e) Establish "no silent updates" governance for material AI dependencies (vendor model/API upgrades, retrieval corpus changes, routing/guardrail changes, permission changes); and (f) Begin internal audit/second-line testing of inventory completeness and control operation for the highest tier. Wave 3 (Months 6–12): Demonstrate operating effectiveness and resilience (a) Complete independent validation/testing for high-impact AI use cases and material GenAI/agent deployments as applicable; (b) Demonstrate monitoring effectiveness (drift detection, anomaly triggers, escalation) and correction closure for material AI; (c) Conduct containment and recovery drills for high-risk AI, including vendor outage/correlated update scenarios and tool-permission containment;

		(d) Mature concentration/coupling monitoring for shared model providers/endpoints and correlated update risk; (e) Establish evidence readiness such that key artefacts can be produced within supervisory request windows (inventory extracts, validation records, change logs, incidents, exceptions, and remediation status); and (f) Complete internal audit coverage on a risk-based cycle for high-risk AI and confirm remediation of priority gaps. 3. Proportionality during transition (avoid undue burden). 3.1 MAS may clarify that smaller or lower-AI-exposure FIs can adopt simplified governance structures (e.g., combined committees or managed validation/testing) provided baseline obligations are met and higher-risk use cases receive appropriate controls. 3.2 MAS may also clarify that low-impact AI can transition primarily through identification, inventory, baseline controls, and reasonable monitoring, while high-impact AI must demonstrate operating effectiveness by the end of the transition. 4. Recommended supervisory emphasis (what to prioritise in the first year). 4.1 The highest supervisory value in the first 12 months is achieved by prioritising: (a) completeness and accuracy of AI identification and inventory (including embedded/vendor and staff-facing AI); (b) consistent tiering/materiality assessment; (c) operating effectiveness of change controls (including vendor update governance); (d) incident readiness and containment capabilities; and (e) third-party risk controls where reliance is material. 5. Optional enhancement: staged compliance declaration (clarity and accountability). 5.1 MAS may consider allowing (or requesting) staged attestations at the end of each wave (0–3, 3–6, 6–12 months), focused on the exit criteria above, to promote disciplined execution and reduce end-of-period compression risk. 6. Summary. 6.1 A 12-month transition is supported and is practical if structured into phased implementation waves with clear exit criteria. 6.2 Early emphasis should be on identification, inventory, tiering and high-risk controls; later emphasis on operating effectiveness, resilience drills, third-party governance maturity, and evidence readiness. 6.3 This approach preserves proportionality while ensuring that the most consequential AI deployments are fully governed and controlled within the transition period.
38	Global Legal Entity Identifier Foundation	**Question 1:** **Question 2:** **Question 3:** **Question 4:** **Question 5:**

		Question 6: **Question 7:** **Question 8:** The Global Legal Entity Identifier Foundation (GLEIF) welcomes the Monetary Authority of Singapore's (MAS) leadership in progressing a future-proof framework for fund tokenisation. GLEIF welcomes MAS's efforts to further strengthen the Singapore's artificial intelligence (AI) framework. GLEIF would like to respond to 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 financial institutions (Fis) should assess for relevance to the AI model, system or use case, and apply in a proportionate manner. GLEIF recommends that FIs recognize that LEI (Legal Entity Identifier, ISO 17442) as a data set that can support AI frameworks for the purposes of identification and reference data of organizations. The LEI system offers a consistent and reliable dataset of legal entities, which is crucial for training and testing AI models to ensure they produce accurate and reliable results across different jurisdictions. The LEI can also be used to identify the entity that owns or operates the AI model. The LEI repository provides high-quality, standardized, open-source legal entity data. These are prerequisites for any good data analysis project or an AI model. Relying on global standards does not just ensure consistency, it increases the data quality and offers a ready-to-use labeled data set for developing Machine Learning (ML) and AI models. Benefits of the LEI: • Improves entity identification: AI models can use the LEI to accurately identify and distinguish between entities, even those with similar names. The LEI is linked to key reference data like the official name, registered address, and country of formation, which helps in identifying entities for due diligence and verification. • Enhances fraud detection and risk assessment: By linking LEIs with other datasets, like beneficial ownership and sanctions data, AI can create more effective systems for detecting financial crime and assessing risk. • Automates compliance and due diligence: AI can use LEIs to significantly speed up processes that require entity identification, such as know-your-customer (KYC) and anti-money laundering (AML) checks, by automating due diligence and compliance tasks. • Simplifies complex data analysis: The LEI system helps AI models analyze complex ownership structures. For example, by integrating with Large Language Models, AI can extract and analyze information from unstructured documents to map ownership structures more accurately. GLEIF remains at MAS' disposal to further discuss and support your work as well as to engage in discussions and your questions related to standards. **Question 9:**

		Question 10:
39	Grant S Case	**Question 1:** **Question 2:** **Question 3:** **Question 4:** **Question 5:** **Question 6:** **Question 7:** Feedback on MAS AI Materiality Dimensions MAS proposes three dimensions for assessing AI materiality: impact, complexity, and reliance. I broadly agree with this structure and focus my comments on the reliance dimension, where additional implementation guidance may support more consistent application. Current Assessment Landscape Impact and complexity are already assessed through established institutional processes. Impact aligns with existing consequence and materiality assessments, while complexity is addressed through model risk management and validation practices. Reliance, as described, does not yet map to an existing supervisory construct, though operational continuity frameworks offer a natural foundation. Paragraph 3.10(c) identifies autonomy, degree of human oversight, and availability of alternatives as relevant considerations. These factors are appropriate. However, in the absence of guidance on how they should be assessed or evidenced, institutions are likely to apply them differently. In particular, the availability of alternatives may be interpreted narrowly, as theoretical substitutability, rather than operational feasibility under normal control and volume conditions. Integration with Existing Frameworks In practice, financial institutions already assess operational reliance through business continuity and resilience frameworks. Systems included in business continuity plans, and assigned recovery time objectives, reflect an established internal view of operational criticality. Where AI systems are embedded in these processes, the associated documentation provides observable and auditable indicators of reliance. I therefore suggest that reliance be assessed in terms of operational continuity duration: the period for which a business function can continue to deliver expected outcomes following the unavailability of an AI system.

		Benefits of Time-Based Assessment This framing provides a measurable basis for reliance assessment and supports comparability across institutions. It also aligns with the Annex consideration of whether loss of access to AI services would disrupt workflows on which the institution is materially dependent. Expressing reliance in time-based terms allows proportionality to be calibrated more clearly, distinguishing between systems with immediate operational dependency and those with longer viable workarounds. From an implementation perspective, shorter disruption tolerances may warrant greater supervisory attention than longer ones. The operational difference between disruption within one day and within several days is not linear, and institutions already reflect this distinction in resilience and continuity planning. Holistic Implementation Reliance assessment should be considered alongside, rather than in place of, impact and complexity. A system may have low operational reliance but high impact when it fails, or low reliance but elevated model risk due to opacity. All three dimensions remain relevant. My observation is that reliance would benefit from clearer measurement guidance to support consistent supervisory outcomes. I would be supportive of further engagement on how reliance assessment could be operationalised using existing resilience and continuity frameworks. **Question 8:** **Question 9:** **Question 10:**
40	Holland & Marie Pte Ltd	**Question 1:** We agree with the principle of applying the Guidelines in a proportionate manner. In this context, we believe it is critical to consider the dual nature of AI when assessing proportionality. The Guidelines rightly identify that the unique risks of AI necessitate more intensive controls, which will add a significant implementation effort. However, we believe it is critical to view this burden in the context of AI's dual nature: it is both a source of risk and a powerful risk mitigation tool. For instance, AI can provide substantially broader coverage of communications than traditional sampling-based controls, or automate aspects of transaction monitoring to enhance detection effectiveness. This perspective helps contextualize the investment in new controls. The principle of proportionality, central to the Guidelines, is the mechanism that allows FIs to manage this balance. We respectfully propose that the Guidelines explicitly acknowledge not only the new risks AI creates, but also the efficiencies it may provide. The

		Guidelines could state that, as part of a proportionate approach, the successful implementation of robust AI-driven controls can justify the modification or retirement of legacy controls they are designed to replace. Such control substitution would be subject to rigorous internal governance—including evidence-based validation, ongoing performance monitoring, and defined fallback plans—to ensure the overall control environment is demonstrably more effective. **Question 2:** **Question 3:** **Question 4:** We offer the following comments on the proposal for FIs to establish a dedicated cross-functional committee to oversee AI risk if the overall AI risk exposure is deemed material. Many of the principles in the Guidelines mirror the sound governance already expected within FIs' existing enterprise risk management frameworks. We strongly agree that the primary value of the Guidelines is to buttress these existing frameworks with a precise focus on the unique risks of AI. To that end, we respectfully submit that the Guidelines could more strongly emphasize the integration of AI risk management into existing structures. Language in paragraph 2.4 of the Guidelines, which suggests establishing a "dedicated cross-functional committee" where the FI's overall AI risk exposure is deemed material, could be interpreted as implying a default expectation to create a new governance body, which may inadvertently lead to duplicative structures or siloed oversight. We respectfully propose that the Authority clarify that an FI can meet the expectation for robust, cross-functional oversight by leveraging and enhancing existing structures. Specifically, the final Guidelines could state that this oversight may be satisfied by expanding the mandate of an existing committee or by implementing a federated governance model with formalized cross-functional coordination. This would support consistent standards and clear accountability across the FI, while avoiding unnecessary duplication and preserving the Authority's supervisory goals. **Question 5:** **Question 6:** **Question 7:** **Question 8:** **Question 9:** In response to the invitation for comments on other aspects of the Guidelines, we wish to highlight the importance of explicitly integrating the forthcoming Project MindForge handbook to ensure effective implementation.

		We welcome the reference in paragraph 2.5 of the Consultation Paper to the industry-led Project MindForge handbook. To maximize its utility, the final Guidelines should contain explicit, in-line references directing FIs to this practical guidance. FIs will look to the Guidelines as the primary source of regulatory expectations, and without such signposting, they may overlook the invaluable, context-specific implementation details the handbook will provide. We respectfully propose that the Authority embed these references within the final Guidelines and potentially include a dedicated "Companion Guidance" section or box that directs FIs to key practical supplements, such as the Project MindForge handbook (once issued) and other relevant information papers or designated toolkits. This would create a clear, multi-layered framework and ensure FIs can effectively translate principles into practice. These comments are offered in support of the Authority's objective to foster trustworthy and effective AI adoption in the financial sector. We once again thank the Authority for its leadership in this critical area and for its consideration of our comments. We would be pleased to provide any further clarification required. **Question 10:**
41	Interactive Brokers Singapore Pte Ltd	**Question 1:** No comments **Question 2:** No comments **Question 3:** No comments **Question 4:** FIs face and manage different types of risks, including risks arising from use of AI. We consider that there is no need to specify a specific forum or a dedicated committee for AI-related issues to be discussed, as long as there are commensurate governance and controls in place. **Question 5:** It would be useful for MAS to provide examples (e.g. what are types of criteria or use cases). **Question 6:** No comments **Question 7:** No comments **Question 8:** No comments

		Question 9: No comments **Question 10:** No comments
42	International Organization for Standardization (ISO), Standards Advisory Group (SAG)	**Question 1:** **Question 2:** **Question 3:** **Question 4:** **Question 5:** **Question 6:** **Question 7:** **Question 8:** The SAG welcomes MAS's efforts to further strengthen the Singapore's artificial intelligence (AI) framework. The SAG would like to respond to 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 financial institutions (Fis) should assess for relevance to the AI model, system or use case, and apply in a proportionate manner. The SAG recommends that FIs recognize that LEI (Legal Entity Identifier, ISO 17442) as a data set that can support AI frameworks for the purposes of identification and reference data of organizations. The LEI system offers a consistent and reliable dataset of legal entities, which is crucial for training and testing AI models to ensure they produce accurate and reliable results across different jurisdictions. The LEI can also be used to identify the entity that owns or operates the AI model. The LEI repository provides high-quality, standardized, open-source legal entity data. These are prerequisites for any good data analysis project or an AI model. Relying on global standards does not just ensure consistency, it increases the data quality and offers a ready-to-use labeled data set for developing Machine Learning (ML) and AI models. Benefits of the LEI: • Improves entity identification: AI models can use the LEI to accurately identify and distinguish between entities, even those with similar names. The LEI is linked to key reference data like the official name, registered address, and country of formation, which helps in identifying entities for due diligence and verification. • Enhances fraud detection and risk assessment: By linking LEIs with other datasets, like beneficial ownership and sanctions data, AI can create more effective systems for detecting financial crime and assessing risk.

		• Automates compliance and due diligence: AI can use LEIs to significantly speed up processes that require entity identification, such as know-your-customer (KYC) and anti-money laundering (AML) checks, by automating due diligence and compliance tasks. • Simplifies complex data analysis: The LEI system helps AI models analyze complex ownership structures. For example, by integrating with Large Language Models, AI can extract and analyze information from unstructured documents to map ownership structures more accurately. The SAG remains at MAS' disposal to further discuss and support your work as well as to engage in discussions and your questions related to standards. **Question 9:** **Question 10:**
43	Investment Company Institute	**Question 1:** The Investment Company Institute (ICI) [1] appreciates the opportunity to provide feedback on the Monetary Authority of Singapore (MAS) consultation on the guidelines on Artificial Intelligence (AI) risk management ("Consultation") [2]. ICI's mission is to strengthen the foundation of the asset management industry for the ultimate benefit of the long-term individual investor. Our members have significant experience and expertise in serving more than 120 million retail investors with US$54.3 trillion in assets under management in jurisdictions around the world, including in the United States, Singapore, Japan, Australia, the United Kingdom, and the European Union. These investors rely on regulated investment funds to achieve important financial goals, including saving for retirement and education. ICI has a strong interest in ensuring that regulatory frameworks encourage innovation while preserving effective standards of liquidity, transparency, and investor protection. We believe that technology has tremendous potential to benefit investors, and that is critical for global regulators to proceed deliberatively, with the goal of preserving existing policies that foster liquid and efficient capital markets while facilitating further innovation. ICI welcomes MAS' work in collaboration with the financial industry to provide thought leadership and to develop good practices with regard to the use of AI. Nevertheless, we are concerned about MAS' approach outlined in this consultation to introduce detailed and relatively prescriptive regulatory guidelines specifically for the use of AI. In our view, existing, technology-agnostic rules and regulations, which incorporate well-established legal principles and standards, are adaptable enough to address fund managers' use of AI and other emerging technologies. Fund managers are already subject to extensive regulatory obligations, for example, to provide true, accurate, and fair information, and to maintain robust cybersecurity and operational resilience, regardless of whether AI is adopted. Rather than creating new regulatory guidelines, we encourage MAS to focus on work such as the AI risk management resources being developed with industry through MAS' Project MindForge. This type of collaborative work with industry can provide fund managers with useful references and good practices as they adopt AI and

		other emerging technologies in a manner consistent with existing regulatory obligations and to benefit the retail investors that they serve. In contrast, a prescriptive approach specifically for uses of certain emerging technologies such as AI will risk imposing a chilling effect on fund managers' willingness to engage in technological innovation with respect to investor services and products, and could deter fund managers from adopting beneficial AI tools in their operations. Further, the International Organization of Securities Commissions (IOSCO) is in the process of evaluating whether it is appropriate to develop additional tools, recommendations, or considerations for the use of AI, taking into consideration existing regulatory frameworks and risk management practices. To enhance international regulatory interoperability, we encourage MAS to take stock of the results of these discussions before moving forward with its own regulatory guidelines. If MAS determines to move ahead with the proposed regulatory guidelines on AI risk management for financial institutions, it is important that such guidance be applied in a proportionate as well as flexible manner. The approach should provide fund managers sufficient flexibility in the methodology for assessing the materiality of AI risks and in selecting an appropriate framework to manage these risks. Such flexibility should include the ability to adapt existing governance, oversight bodies, and risk management processes, or to use tailored controls and tools, provided that these controls can provide adequate oversight. We also recommend that any guidance on AI risk management should allow global fund managers to leverage group resources, staff, and policies and procedures to satisfy the MAS regulatory expectations regarding AI governance and risk management. Thank you again for the opportunity to provide feedback on the Consultation. As AI continues to evolve, we look forward to engaging with the MAS on AI policies that promote innovation and efficiency and protect our capital markets. Note: [1] The Investment Company Institute (ICI) is the leading association representing the global asset management industry in service of individual investors. ICI members are located in North America, Europe, and Asia and manage fund assets of US$54.3 trillion, including mutual funds, exchange-traded funds (ETFs), UCITS, closed-end funds, unit investment trusts (UITs) and similar funds in these different jurisdictions. ICI has offices in Washington, DC, Brussels, and London. [2] Monetary Authority of Singapore, Consultation Paper on Proposed Guidelines on Artificial Intelligence Risk Management for Financial Institutions (13 November 2025), *https://www.mas.gov.sg/publications/consultations/2025/consultation-paper-on-guidelines-on-artificial-intelligence-risk-management*. **Question 2:** **Question 3:**

		Question 4: **Question 5:** **Question 6:** **Question 7:** **Question 8:** **Question 9:** **Question 10:**
44	Investment Management Association of Singapore	**Question 1:** IMAS members welcome the introduction of the principle of proportionate application and encourage further clarification to ensure the Guidelines are risk-focused and operationally workable while maintaining robust oversight for AI that materially affects decisions, customers, or markets. While the policy intent is clear, that said, the industry is still concerned about the following: (a) Risk of Over-Scoping Low-Risk AI Uses Our members expressed concern that the current framing of AI "integrated into business processes" may unintentionally capture low-risk, ancillary AI tools that do not materially impact decisions, customers, or markets. Examples include AI used for translation, summarisation, document search, productivity enhancement, or standard contract review. Although technically embedded in workflows, these tools do not exercise discretion or substitute human accountability; do not materially influence outcomes nor are they subject to human review or are easily reversible. Applying the full suite of governance, documentation, and oversight requirements to such use cases could create disproportionate compliance burdens and discourage responsible adoption of benign AI. MAS is encouraged to clarify that AI tools which are supportive or administrative in nature, not decision-making or discretion-exercising, and not subject to human oversight may justifiably be subject to 'lighter-touch requirements' and clarify what such 'lighter-touch requirements' might be, under the proportionality principle even if they are embedded within business processes. (b) How proportionality should be applied Our members noted that further clarity is needed to support consistent and practical implementation, including: • How to distinguish between basic AI usage and AI use cases requiring enhanced controls,

		• Whether all FIs are expected to maintain baseline AI policies even in the absence of AI use cases, and • How governance, monitoring, and assurance expectations should scale with AI risk levels. Hence, we suggest that the Guidelines should explicitly state that proportional application is primarily determined by whether an AI application constitutes a formal AI use case and/or materially impacts, informs, or supports decisions for which the FI is accountable. This would help focus governance and resources on AI systems that present meaningful risk, while avoiding unnecessary controls for non-material AI usage. To support implementation, MAS could provide illustrative examples of AI risk tiers (e.g. low, moderate, high risk) and demonstrate how governance, monitoring, and assurance expectations scale accordingly. (c) Operational flexibility for multinational FIs AI governance There is concern that requirements such as designating a senior management member responsible for AI oversight may not sufficiently reflect global operating models, particularly for FIs operating across multiple jurisdictions. While accountability should remain with the local FI, rigid expectations could be operationally challenging where AI expertise and oversight are centralised at the parent or intragroup level. We suggest that MAS clarifies that while ultimate accountability rests with the local FI, the designated senior management member may leverage parent or intragroup AI expertise and frameworks to discharge AI oversight responsibilities, consistent with existing guidance on group-level reliance. An industry engagement session by the MAS to clarify all of the above would be helpful. **Question 2:** The comments from members are: (a) Overly broad definition of AI risks capturing low-risk, deterministic tools There is a shared concern that the proposed scope may unintentionally capture traditional analytical and automation tools that are deterministic, transparent, and non-adaptive. These tools—such as fixed-parameter regression models, deterministic statistical models, scoring formulas, spreadsheet macros, rule engines, and non-learning RPA—have long been used in regulated environments and do not exhibit learning, evolution, or emergent behaviour. Subjecting such tools to the full AI governance framework could result in disproportionate regulatory burden without commensurate risk mitigation benefits. MAS is therefore encouraged to clearly state that purely deterministic, non-learning tools—including fixed-parameter regression models, macros, rule engines, deterministic statistical models, and non-learning RPA—are not regarded as AI use cases, systems, or models under the Guidelines. This would prevent the inadvertent inclusion of traditional, low-risk tools with predictable behaviour. (b) Insufficient Clarity Around What Constitutes an “AI Use Case”

		Respondents emphasise that the definition of an AI "use case" is critical to ensuring a risk-based and proportionate application of the Guidelines. With the increasing integration of AI-augmented features in standard third-party software, there is a risk that entire platforms or software suites could be classified as AI use cases, even where AI outputs are merely assistive and do not affect formal decision-making or accountability. We suggest that MAS refines the definition of an AI use case to anchor on material risk and accountability, to focus on situations where AI outputs are integrated into business processes in a way that formally impacts accountabilities, responsibilities, or decision-making of the FI. Where AI does not substitute or affect human accountability, its use should be treated as non-material risk and fall outside the core scope of the Guidelines. (c) Need to Distinguish Decision-Critical AI from Assistive or Productivity Tools A key concern is the lack of clear differentiation between AI that formally impacts or substitutes human decision-making and accountability, and AI that is purely assistive or productivity-enhancing (e.g. drafting, summarisation, "help" chatbots). Where human accountability is not substituted and AI outputs do not drive decisions, such uses pose non-material risk and should not fall within the substantive scope of the Guidelines. It would be helpful if MAS provides explicit clarification that commonly used productivity tools (e.g. AI-assisted drafting, summarisation, analysis features) and embedded "help" or chatbot functionalities do not constitute AI use cases under the Guidelines, provided they do not influence formal decisions or accountability. (d) Practical Challenges in Managing Third-Party Embedded AI While respondents agree with the inclusion of third-party AI within scope, concerns were raised about practical implementation where AI is deeply embedded in vendor platforms and cannot be selectively disabled or modified by the FI. Without additional guidance, FIs may face uncertainty in meeting regulatory expectations despite limited control over the AI components. To address operational realities, MAS could provide examples of compensatory controls, risk acceptance frameworks, or governance approaches applicable where AI is embedded within third-party platforms and cannot be disabled. This would support consistent and pragmatic implementation across FIs. **Question 3:** Members generally support strong Board and senior management accountability for AI risks, but advocate for a proportionate, flexible, and principle-based approach. To sustain clear oversight and mitigate governance and control fatigue, members suggest that these responsibilities should be applied proportionately on AI use cases only in circumstances where the threshold set out in Question 2 is met (i.e. output, whether from an internal system or a feature within a third-party product, which is integrated into a business process in a manner that formally impacts accountabilities, responsibilities, and decision-making of the FI). The FI's Board's primary oversight should be mandated for AI that is integrated into accountable business processes and not diffused across all AI-assisted activities. This focuses Board and senior management's accountability where human leadership attention matters most.

		Some members highlighted that as a part of global fund management companies, their companies are currently subject to existing laws e.g. the EU AI Act which European-headquartered FIs have to comply with and are in the process of implementing group-wide AI governance frameworks. The EU requirements are extensive. In practice: • AI oversight and expertise may be distributed across group level governance structures. • AI risk identification and classification (e.g. high-risk determination, prohibited use cases) are performed centrally, with local entities relying on group-approved frameworks and tooling. • AI risk assessments may involve cross functional processes across business, compliance, legal, risk, information security, IT and third-party risk. Creating additional entity level committees or assigning all responsibilities to one function may lead to duplication, operational complexity and misalignment with existing governance models. It is also not very clear as to how the Board and Senior Management may divide and discharge their duties and responsibilities if a cross-functional AI committee is established. It is recommended that MAS adopt a principle based and flexible approach that acknowledges how multinational FIs organise AI governance. Specifically: • Allow FIs to rely on group level governance and distributed oversight models (similar to those under the MAS Outsourcing Guidelines), rather than concentrating accountability on a single local individual or the local Board and/or senior management. • Avoid mandating a new, separate AI oversight committee for material risk AI use case where existing cross functional forums already provide adequate challenge and risk management. • Permit responsibilities for AI identification, inventory maintenance and risk assessment to be distributed across the relevant specialist functions, instead of requiring a single designated control function. This approach preserves accountability while giving FIs the flexibility to leverage (and learn from) established frameworks, avoid duplicative governance, and organise responsibilities where the expertise resides. To better guide the industry, members suggest MAS consider providing illustrative examples of ongoing monitoring approaches for AI risks, including examples of measures, limits or key risk indicators that FIs may consider, calibrated to the risk materiality of different AI use cases. Such practical examples would help FIs translate governance expectations into operational monitoring and Board and governance reporting practices, particularly in areas such as third-party AI, customer-impacting AI use cases, and AI embedded within critical systems. **Question 4:** IMAS respectfully suggests that the Guidelines should require effective and empowered oversight of material AI risks by FIs and not prescribe a specific committee structure. FIs should be allowed the flexibility to integrate this oversight into existing governance committees or equivalents that already bear ultimate accountability, provided they can

		demonstrate that the forum has the competent authorities. This reinforces and empowers established accountability structures and mitigates fragmentation. Equally, the FI could decide to establish a dedicated cross-functional oversight committee for AI risk management coverage if it deems fit. It is also suggested that MAS recognises that cross-functional oversight, including AI expertise and associated control functions, could be drawn from capabilities within a parent / intragroup entity which facilitates substantive oversight whilst benefiting from economies of scale arising from leveraging established group-level resources. **Question 5:** Our members view systematic identification as being critical for managing risks associated with AI use cases where the threshold set out in Question 2 is met (i.e. output, whether from an internal system or a feature within a third-party product, which is integrated into a business process in a manner that formally impacts accountabilities, responsibilities, and decision-making of the FI). We would suggest that the expectation for systematic identification should be clearly calibrated to the aforementioned AI use cases. Paragraph 3.3 of the Guidelines sets out that "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". FMCs with regional/global operations find it beneficial to draw on AI expertise and associated control functions (or equivalent) from a parent / intragroup entity for AI identification systems and processes. We would also seek confirmation from MAS that it is acceptable for the AI risk management framework and policy to be facilitated by a primary function, while specific components of the framework are owned and implemented by different relevant stakeholders, provided that roles, responsibilities and accountability are clearly defined. For example, technical and due diligence may be owned by IT, high-level definitions and processes by risk management and compliance, and contractual standards by legal. Such a federated operating model would support effective implementation while maintaining clear accountability and governance. **Question 6:** IMAS members expressed the following concerns: (a) Risk of overly broad inventory scope A requirement to inventory all AI usage could unintentionally capture low-risk, assistive, or ancillary AI tools that do not materially impact decision-making or accountability. Applying inventory requirements indiscriminately risks creating unnecessary operational burden and diluting focus on AI use cases that present material risk. (b) Proportionality and risk-based calibration While members generally support the establishment of a comprehensive AI inventory, there is a need for flexibility in the depth and granularity of information captured. A uniform level of detail across all AI uses may not be proportionate, particularly for lower-risk AI applications.

		(c) Organisational practicalities Members which are part of multinational FIs, AI governance, expertise, and tooling are often centralised at the parent or group level. Requiring inventory ownership and execution solely at the local entity level may result in duplication and misalignment with existing group-wide governance models. (d) Clarity on inventory attributes and integration with existing registers A concern was raised about the lack of clarity around certain proposed inventory attributes, particularly "model type", and how this information is expected to be used for governance and risk assessment. There is also uncertainty around how the AI inventory should link to other existing inventories, such as system asset inventories and third-party or outsourcing registers. To address the above, we respectfully suggest MAS considers the following possible solutions: (a) Limit inventory requirements to material AI use cases MAS is encouraged to clarify that AI inventory requirements should apply only to AI use cases that meet the materiality threshold, i.e. where AI outputs are integrated into business processes in a manner that formally impacts accountabilities, responsibilities, or decision-making of the FI. This would ensure inventory efforts are focused on AI that poses meaningful risk. (b) Apply proportionate depth and granularity based on risk The Guidelines could explicitly state that while all in-scope AI use cases should be recorded, the level of detail required for each inventory entry may be calibrated according to the assessed risk materiality of the AI use case. Lower-risk AI could be subject to a simplified inventory, with enhanced attributes reserved for higher-risk applications. (c) Permit reliance on parent or intragroup functions MAS should clarify that FIs may leverage parent or intragroup AI expertise, tooling, and control functions to maintain the AI inventory, provided there is clarity on the scope of accountability remaining with the local FI and that with the parent/intragroup company. This would align inventory practices with existing group-wide governance and avoid duplicative efforts. (d) Clarify inventory attributes and their purpose Further guidance on how attributes such as "model type" should be defined and used in AI risk assessment and governance would improve consistency. Illustrative examples would help FIs understand supervisory expectations. (e) Explain expected linkages with other Governance Inventories MAS could elaborate on how the AI inventory should integrate with existing inventories (e.g. system asset registers, outsourcing or third-party registers), to support holistic risk management and reduce duplication across governance frameworks. **Question 7:** Across respondents, there is agreement that the proposed dimensions of impact, complexity, and reliance are appropriate. However, consistent with earlier responses, members caution that applying detailed materiality assessments across all AI usage risks

		over-burdening FIs. The concern is that such assessments should be reserved for AI use cases that formally impact accountabilities, responsibilities, or decision-making (in line with Question 2), where AI use poses material risk to the FI. Respondents noted that while the dimensions are appropriate, the Guidelines provide limited guidance on how to operationalise them consistently. In particular, there is uncertainty over whether MAS expects a standardised scoring framework and how assessments should differ for generative AI or agentic workflows, which may present distinct risk characteristics compared to traditional models. Our members which are part of multinational FIs reiterate concerns around the expectation to assign a single local control function as the final arbiter for risk materiality assessments as AI expertise and assessment frameworks may be centralised at the group level, and strictly localised ownership could create duplication and inefficiencies. We suggest that MAS consider: (a) Gate materiality assessments by decision impact threshold MAS is encouraged to clarify that detailed assessments across the dimensions of impact, complexity, and reliance should apply only where the AI use case meets the materiality threshold, i.e. where AI outputs are integrated into business processes in a way that formally impacts human accountability and decision-making. (b)Anchor the assessment to human accountability outcomes The Guidelines could explicitly state that the assessment should evaluate how each dimension influences the quality, explainability, and contestability of human decisions. It should also ensure alignment with the core governance objective of preserving human accountability. (c) Provide reference scoring frameworks or methodologies To promote consistency, MAS could consider providing a reference scoring framework or illustrative assessment methodology and guidance on whether and how criteria should differ for generative AI and agentic systems. Such guidance would support comparability across FIs while allowing flexibility in implementation. (d)Allow Group-Level Control Functions to Perform Assessments MAS could clarify that while accountability remains with the local FI, risk materiality assessments and supporting control functions may be performed or supported at the parent or intragroup level, leveraging centralised expertise and tools. **Question 8:** We are generally supportive of MAS's comprehensive AI life-cycle framework but call for greater clarity, risk-based calibration, and practical flexibility, particularly for low-risk AI and third-party deployments. However, there are diverse views from the industry on this matter: • Insufficient clarity on proportionate application across different risk levels While members support MAS's end-to-end AI life-cycle framework and the emphasis on proportionality driving risk assessment, there is concern that baseline expectations for

		applying the framework to different AI risk levels are not clearly articulated. Without clearer differentiation, there is a risk that similar oversight, testing, and governance requirements will need to be applied across fundamentally different AI use cases, undermining proportionality and diverting resources away from higher-risk systems. Hence the Guidelines should clarify that the full suite of AI life-cycle controls applies primarily to AI use cases that are material in accordance with the principle of proportionality, i.e. those that formally impact decisions, accountability, or outcomes. Lower-risk or assistive AI should be subject to lighter-touch controls, aligned with their risk profile Illustrative examples or supervisory benchmarks showing how controls (e.g. human oversight, testing, monitoring, assurance) scale across low-, medium-, and high-risk AI systems would support consistent interpretation of proportionality and clearer demonstration of compliance. • Integration with existing outsourcing and technology governance frameworks There is concern that AI life-cycle controls should be implemented as a parallel regime instead of being integrated into existing technology and model risk management frameworks, reducing efficiency and consistency. Members suggest that MAS explicitly encourages FIs to adapt and augment existing technology governance and lifecycle controls (e.g., Software Development Life Cycle, change management, monitoring) to address AI-specific risks, rather than requiring entirely new parallel frameworks. • Governance Our members seek clarification on supervisory expectations where, after applying reasonable mitigants, certain risks cannot be fully assessed or addressed. In particular, guidance on whether such residual risks should be explicitly escalated for senior management or Board risk acceptance. • Industry training / capacity building. MAS could consider promoting or facilitating industry-level training and capacity-building initiatives on AI-specific risk areas such as fairness and explainability, to enhance collective expertise and support consistent implementation Where residual risks remain after reasonable mitigants, there is clear guidance on escalation, senior management or Board risk acceptance, and documentation expectations. Specific points for Section 6 Proposed Guidelines on AI Risk Management • Para 3.8 to 3.11 AI Risk Materiality Assessment It would be helpful if MAS can provide illustrative examples or supervisory benchmarks for low-, medium-, and high-risk AI systems to enhance clarity and support consistent application of the Guidelines. This would allow FIs to calibrate controls appropriately, focus enhanced oversight where it is most needed, and demonstrate compliance in a transparent manner.

		Further, for many AI systems, particularly third-party or foundation models, FIs act as deployers rather than developers or providers. In these cases, core aspects of model governance (including model design, training, tuning, validation, and underlying performance testing) sit squarely with the provider and are outside the direct control of the FI. • Para 4.3 AI Life Cycle Controls A common concern is that FIs often act as deployers rather than developers of AI, especially for third-party or foundation models. In such cases: model design, training, validation, and detailed testing are outside the FI's direct control. Visibility into fairness, transparency, and supply chain risks may be limited and usually, model details are treated as proprietary intellectual property and the third parties will not share the models. This raises uncertainty around how FIs can meet life-cycle control expectations in a meaningful and proportionate way. Our members suggest that MAS clarifies if reliance on vendor documentation, independent assessments (if available), and periodic external evaluations (if available) is acceptable where direct access to model details is not available. Compensatory controls (beyond limiting use) may be applied where full transparency or fairness testing is impractical. Firm-level versus use-case-level expectations for concentration risk assessments and contingency planning should be clarified, including whether qualitative assessments are sufficient. • Para 4.6 Transparency and Explainability MAS is requested to provide guidance on the level of disclosures and the standards of explainability. Please provide an illustration on transparency, explainability and the relevant controls. • Para 4.8 / 4.9 Fairness Members seek clarification that fairness testing is not required for GenAI systems in circumstances where the threshold set out in Question 2 is not met (i.e. AI output, whether from an internal system or a feature within a third-party product, is not integrated into a business process in a manner that formally impacts accountabilities, responsibilities, and decision-making of the FI) as this is not where use of AI poses material risk. • Para 4.10 Human Oversight While members generally understand and accept the principle of ensuring appropriate human oversight and apply proportionate evaluation and testing, with enhanced controls commensurate with risk, they seek clarification on the baseline expectations for different risk tiers. This is as applying similar oversight and testing requirements across fundamentally different use cases risks undermines the proportionality principle and diverts resources away from genuinely higher-risk systems. IMAS members seek further clarification on the minimum standards MAS expects for human oversight and testing, particularly for lower risk AI systems where enhanced controls may not be necessary. Clearer guidance would help FIs apply proportionality consistently while maintaining risk appropriate governance. • Para 4.11 Third-Party AI Management

		The proposed Guidelines expect an FI to ensure that onboarding, development and deployment controls for third-party AI are adequate for the risk materiality of the use case, system or model which uses or depends on third-party AI. Members seek MAS' advise on how Fis should meet the standard on transparency, fairness, supply chain assessments, concentration risk, contingency plans etc. • Para 4.11 Concentration Risk Members seek clarification from MAS if the assessment should be from the corporate perspective e.g over-reliance on key third-party AI providers (whether direct or indirect) on the firm level or on application level? Please provide some clarification on how to define "over-reliance", is qualitative analysis sufficient or some quantitative figures should also be considered? • Para 4.11 Contingency Plans MAS is requested to provide clarification on whether the contingency plan should be on firm level or expected on each use case where the risk materiality is high. • Paras 4.12 and 4.13 Selection As FIs will generally adopt few AI models as firm standards, there will not be a list of options to select for each application / use case. We would be grateful if MAS can advise its position on firmwide AI model selection - whether this can fulfil the requirement. **Question 9:** IMAS members are of the opinion that the proposed Guidelines provide a comprehensive foundation for managing material AI risks. The following are suggestions for MAS' further consideration in order to enhance clarity and practical applicability of the Guidelines. (a) Clarity on Treatment of Embedded Third-Party AI Members suggested clear guidance on how AI functionality embedded within broader third-party software should be assessed. Without this clarity, FIs may struggle to determine whether such AI features trigger the full application of the AIRG framework, particularly for assistive or productivity-focused functions. There is a risk of either over-applying controls or under-managing material AI risks. Nevertheless, MAS could provide explicit guidance that the AIRG framework should apply only when the AI feature meets the materiality threshold (output integrated into a business process affecting accountabilities or decision-making). Assistive or non-material AI functionality can be managed through the FI's existing third-party vendor and technology risk frameworks. (b) Need for an Overarching Governance Principle IMAS and its members are of the opinion that there should be a clear, principle-based anchor framing the purpose of AI Risk Management. Without this, there is a risk of inconsistent interpretation of proportionality, governance expectations, and the role of AI in supporting effective human decision-making. A member suggested the following wordings: "The purpose of AI risk management is to ensure that the use of AI supports and does not undermine the FI's capacity for effective human judgment, accountability, and its ability to meet its regulatory and conduct obligations.".

		(c) Affirming a Principles-Based Approach to Competency Requirements While members support the Guidelines' emphasis on competency and outcomes over prescriptive roles, it suggests explicitly reaffirming that compliance should focus on skills, governance capabilities, and effective oversight rather than specific titles or functions. This would promote flexible, practical implementation across FIs with different organisational structures. MAS could explicitly emphasise that FIs should ensure the presence of necessary skills and effective governance processes, regardless of whether these are provided internally, via external hires, or through managed services. This supports flexible compliance pathways while maintaining robust oversight. (d) Practical Guidance for Third-Party AI service providers and Assurance Limitations One member noted challenges with third-party AI service providers, particularly where independent assurance reports (like OSPAR or ISAE) are not available. FIs may need guidance on what constitutes acceptable alternative assurance measures or compensatory controls, especially when managing embedded AI in outsourcing arrangements. **Question 10:** IMAS members generally agree with the proposed transition period of 12 months although some have suggested that for firms that are just starting their AI journey, a transition period of 18 months will be appreciated. This is as more time would be required for developing appropriate governance structures, processes, infrastructure and competence from the ground up, particularly for FIs that are only in the nascent stage of adopting AI use cases, systems and/or models.
45	James Lee	**Question 1:** Could the definition and guidance on "proportionate to the size and nature of the FI's activities" be clearer? For instance, smaller FIs may conclude that they are automatically exempt from stricter AI controls even if they adopt high-impact credit scoring AI, autonomous underwriting platforms, AML/CFT Ai engines or biometric onboarding with deepfake risk. Also, nature may also be misinterpreted. For instance, retail lending vs corporate advisory or, B2B instead of retail, hence fairness controls don't apply, or explain-ability is less important to investment banking. FI may also argue that they are not AI-heavy institution, therefore controls are minimal, and this creates shadow AI, where embedded AI in third party products may also avoid oversight. Last but not least, proportionate may also be interpreted as AI maturity level in FI activities. For instance, early adoption may impose simpler controls and bypass foundational controls such as dataset governance, model cards, OWASP LLM Top 10 adversarial testing or human feedback in the model adoption. **Question 2:** While the broad definition covers ML, Deep Learning, Generative AI, and autonomous agents is appropriate and aligned with the EU AI Act, AWS AI Governance Pillars, and ISO/IEC 42001's organisation-wide scope. Could the guidelines also include the following: - AI embedded in vendor systems (e.g., CRM scoring, cloud threat detection) - LLM-based agents capable of autonomous tool/API actions - AI-enabled fraud or sanctions screening components

		This is because many FI centric tech vendors have already embed AI into their solutions. For example, netapp storage had AI that auto categorize data placed in the SAN appliances. SIEM tools, too had AI agents to provide triage advisory. This list could grow larger as more realize the productivity benefits. **Question 3:** In addition to the clear delineation of Board and Senior Management responsibilities, perhaps, Boards should also ensure the following and not leave these to senior management or governance working forums to decide: - Defined AI risk appetite (bias tolerance, transparency expectations, model autonomy thresholds). - Oversight of ethics, fairness, and customer harm mitigation. - Monitoring of AI-linked operational resilience and cyber threats. **Question 4:** Committee membership should include the following functions and a minimum quorum should also be established for major decision on AI usage: Risk, Compliance, IT, Cybersecurity, Data Science, Legal, Customer Experience, and Internal Audit (as independent observer). Presence of Line 3 is important here, as an independent observer; IA could assess the committee's effectiveness and independence, challenge robustness or escalate potential issues. **Question 5:** Perhaps this section can draw reference to international frameworks such as: -learning/inference capability (ISO 42001) - decision influence (EU AI Act) - attack surface exposure (OWASP LLM01–LLM10) For consideration on dataset sensitivity and autonomy level, it can be based on FI's data or information classification. **Question 6:** Can MAS clarify on the expectations of inventory maintenance? In my personal opinion, it should minimally cover these: - model lineage, datasets, training sources - third-party model dependencies - autonomy level (human-in-loop vs human-out-of-loop) - cybersecurity exposure (e.g., prompt injection risk) - drift detection metadata **Question 7:** In addition, data sensitivity and external threat exposure can also be included as these aligned to PDPA, and CSA Securing AI guidelines. **Question 8:** The life cycle approach may be further strengthen to include the following: - Explainability & Transparency - Robustness & Adversarial Security

		- Human oversight **Question 9:** In my personal opinion, I hope MAS could consider the following additional areas: a. For AI-Caused Systemic Outages (Technology Resilience), FI should report as part of the FSM N05 regime. AI-induced outages e.g., automated transaction blocking, routing loops, faulty credit-decision engines, vendor LLM failures should fall under MAS Technology Resilience / TRM incident reporting. b. Requirement for AI Bills of Materials (AI-BOMs) International alignment frameworks (Responsible AI Trust, ISO/IEC) advocate for AI-BOMs to document components, datasets, training sources, and dependencies. c. Change management and monitoring of Prompts Prompt is a new form of query language akin and should be tracked and monitored as part of the AI evolution. For individual who have access to RAG interface of the LLM, prompt should be track or monitored. **Question 10:** NIL
46	Jianjiang Yang	**Question 1:** Drawing on experience across both Global Systemically Important Banks (G-SIBs) and FinTechs, I recommend explicitly distinguishing Major Payment Institutions (MPIs) and Digital Asset Exchanges from Banks in Paragraph 1.5. While Banks require lengthy capital model validation cycles, MPIs operate in high-velocity environments where AI risk management must prioritize System Resilience and Real-time Monitoring over pre-deployment capital adequacy frameworks. **Question 2:** NA. **Question 3:** I recommend modifying Paragraph 2.6(a) to explicitly allow the Board to delegate approval of technical procedures to Senior Management (e.g., Head of Model Risk) or a specialized Risk Committee. To prevent the dilution of strategic oversight, the Board's authority should remain focused on high-level Risk Appetite and Governance Policy. **Question 4:** I recommend allowing FIs to leverage existing governance structures, such as the Model Risk Committee, provided specific AI expertise is represented, rather than mandating a new "dedicated cross-functional committee". Effective oversight is better achieved by integrating AI risk into established risk frameworks supported by a centralized Second Line of Defense (2LOD) rather than creating administrative silos. **Question 5:**

		I recommend explicitly citing Model Risk Management (MRM) as the primary candidate for the "control function" required in Paragraph 3.3. Drawing on industry best practices, the Head of Model Risk should be empowered as the "Final Arbiter" on identification to ensure technical rigor and prevent business lines from misclassifying AI to bypass controls. **Question 6:** I recommend explicitly accepting Automated Model Registries as satisfying the "up-to-date" inventory requirement in Paragraph 3.4. In modern CI/CD environments where models may be retrained weekly, manual document-based inventories are obsolete upon creation; regulations should encourage superior, system-driven governance. **Question 7:** I recommend adding "Reversibility" (or Finality) to the risk dimensions in Paragraph 3.10. In digital assets, an AI agent executing an irreversible on-chain transaction presents a fundamentally higher residual risk than one triggering a reversible fiat transaction. This distinction is critical for accurately assessing materiality in non-banking financial institutions. **Question 8:** I recommend establishing a "Certification Safe Harbor" for Paragraph 4.11. As "Black Box" testing of proprietary vendor models is often impossible without access to weights, FIs should be permitted to rely on independent certifications (e.g., ISO 42001, SOC 2 Type II) to satisfy testing requirements for non-customizable software, rather than attempting technically unfeasible "compensatory testing". **Question 9:** I recommend that MAS explicitly address Open-Source AI Governance as distinct from "Third-Party" vendor relationships. Paragraph 4.11(c) mentions "supply chain assessments" for open-source models, but unlike traditional vendor arrangements, open-source repositories often lack a legal entity to provide "certifications" or "attestations". The Guidelines should explicitly recognize Community-Based Due Diligence (e.g., repository activity, community security audits, and bug bounties) as valid alternative controls for open-source AI. **Question 10:** I recommend a tiered transition timeline: 12 months for internal models and 24 months for Third-Party AI arrangements. Remediation of third-party risks often requires renegotiating global vendor contracts, which is operationally unfeasible within the proposed 12-month window.
47	Kelvin K. F. Law	**Question 1:** Executive summary This response supports MAS's proposed Guidelines with modifications to strengthen implementation. My six principal recommendations are: • Replace the binary "assistive versus integrated" distinction with a decision test plus measurable indicators (Q1). AI use should be classified as "integrated" when output is produced at rates that make independent verification infeasible, when output is incorporated into customer communications without re-performance of key checks, or

		when AI directly triggers actions. A minimum baseline of controls should apply to all AI regardless of classification. • Require an AI Bill of Materials (Q6). The AI inventory should function as a comprehensive bill of materials analogous to software bills of materials (SBOMs) in cybersecurity. It should capture dependencies, versioning, vendor update behaviour, and contractual rights, enabling concentration risk analysis across the AI portfolio. • Publish a standard materiality scoring rubric (Q7). MAS should endorse a simple rubric covering impact, autonomy, scale, reversibility, adversarial exposure, and dependency criticality. A "high" rating on any dimension should trigger enhanced controls, governance, and independent validation. • Specify GenAI and agent-specific controls as explicit expectations (Q8). For generative AI and agentic systems, MAS should expect prompt injection testing, tool permissioning with least-privilege principles, prompt and output logging for audit purposes, and documented kill-switch mechanisms. • Adopt a milestone-based transition with specified deliverables (Q10). Rather than a single 12-month deadline, MAS should require: AI inventory and initial materiality ratings by month 6, controls implemented for high-materiality systems by month 9, and full compliance with independent assurance report at month 12. • Clarify local accountability for offshore AI (Q1, Q2). Where AI materially influences decisions relating to Singapore customers or risks booked in Singapore, the local FI should remain fully responsible for compliance regardless of where the AI is developed, hosted, or accessed. Detailed responses to each consultation question follow. ========== Question 1: Proportionate application Position: Support with modifications. I support the principle of proportionate application. However, the distinction between "assistive" AI and "integrated" AI creates significant implementation challenges that require further guidance. Key concerns • Automation bias risk: The "assistive" carve-out assumes meaningful human review. Research on automation bias (Parasuraman and Riley, 1997; Skitka et al., 1999) and recent studies on LLM overreliance demonstrate that even experts often accept AI-generated errors when outputs appear fluent. A relationship manager reviewing 50+ AI-drafted client letters daily will likely make only cosmetic edits, converting what is nominally "assistive" into de facto "integrated" use.

		• Cosmetic versus substantive review: The Guidelines should clarify what constitutes meaningful human oversight that would qualify AI use as "assistive." Volume-based thresholds or time-per-review minimums may be necessary to prevent cosmetic compliance. • Regulatory arbitrage risk: Business units may deliberately structure AI use to remain within the "assistive" category and avoid governance burdens, even when cumulative effect is material. FIs should be required to assess assistive AI uses in aggregate, not only individually. • Cross-border regulatory arbitrage: In addition to structuring use cases to remain within the assistive category, FIs may also attempt to shift AI development and hosting to offshore entities (or group centres) at the same time they continue to use the outputs to support decisions about Singapore customers or risks. The Guidelines should clarify that where AI materially influences decisions relating to customers or risks booked in Singapore, the local FI remains fully responsible for compliance with the Guidelines, regardless of where the AI is developed or hosted. • Minimum baseline controls for all AI use: To avoid trivializing risk under the assistive label, MAS could specify a simple baseline set of controls that should apply to all AI uses, regardless of whether they are classified as assistive or integrated. At minimum, this might include: (i) registration in an AI inventory; (ii) basic pre-deployment sanity checks (e.g., obvious error, bias, or hallucination tests); and (iii) clear human accountability and a tested ability to disable or roll back the AI tool. • Compliance burden for smaller FIs: I acknowledge the potential for overly broad definitions of "integrated AI" to capture genuinely trivial uses and create disproportionate compliance burdens for smaller FIs such as boutique asset managers or smaller insurers. MAS may wish to consider simplified compliance pathways for FIs below certain thresholds. Illustrative edge case A wealth management firm uses GenAI to draft client portfolio review letters. Each letter is "reviewed" by a relationship manager, but the RM reviews 50+ letters daily and typically makes only cosmetic edits. Is this assistive or integrated? I recommend MAS provide additional guidance on such scenarios. Recommended classification: I recommend MAS treat this scenario as "integrated" AI. The throughput of 50+ letters daily makes substantive independent review implausible. The AI output is the de facto product. Proposed decision rule: AI use should be classified as "integrated" when any of the following conditions applies: (i) output is produced at a rate that makes independent verification infeasible; (ii) output is incorporated into customer communications or decisions without re-performance of key checks; or (iii) output directly triggers an action without human intermediation. AI that does not meet any of these conditions may be classified as "assistive" but should remain subject to baseline controls.

		Question 2: Response Position: Support, with clarifications. I support the broad, technology-neutral definition of AI. However, several emerging architectures and deployment patterns require clarification. Key concerns • Hybrid systems: The Guidelines should explicitly address hybrid systems that combine AI with rule-based components. For example, an AML system may use ML for name standardization before rule-based matching. If the ML component fails, false negatives increase dramatically, yet the ML component is never independently deployed and has no direct customer impact. MAS should clarify that component-level materiality assessments are required even when AI is embedded within larger systems. • Multi-agent orchestration and prompt chains: The consultation's scope definition may struggle with emerging architectures. MAS should clarify that orchestration layers, prompt chains, and retrieval-augmented generation (RAG) pipelines are within scope as AI systems, even if individual components might not meet the "learning/inference" test in isolation. • AI-generated training data: If an FI uses AI to generate synthetic data that trains another model, this creates a lineage problem. Scope should cover such indirect AI dependencies. • APIs with embedded AI: Many traditional vendor APIs (document processing, OCR, translation) now have AI embedded without explicit disclosure. FIs may not even know. I recommend that MAS require FIs to periodically survey vendors on AI content in existing contracts. • Scope independence from provider or location: The Guidelines should clarify that AI use is in scope based on function, not on who develops or hosts the model. Third-party, group-level, and open-source AI models are in scope whenever they materially influence decisions or risk assessments within the FI, even if they are accessed solely via APIs or embedded into vendor products. Illustrative edge case A sanctions screening system uses an ML model to standardize name spellings before rule-based matching. The ML component is never independently deployed and has no direct customer impact, but if it fails, false negatives increase dramatically. How should this be treated under the Guidelines? **Question 3:** Response Position: Support with modifications.

		I support the clear articulation of Board and senior management responsibilities. However, several structural challenges warrant attention. Key concerns • Information asymmetry: There is inherent information asymmetry between technical teams deploying AI and Board members overseeing it. This parallels historical challenges with complex financial products pre-2008. Boards may approve AI deployments without fully understanding their risk profiles. The Guidelines should require FIs to demonstrate that information presented to Boards is comprehensible and enables genuine oversight, not merely formal approval. • Accountability and incentives: Linking accountability to incentives is essential. Senior management compensation structures should reflect AI risk outcomes, not merely deployment volumes or efficiency gains. • Structural conflict between business and risk: Business lines are often better resourced and more influential than risk functions. MAS should expect FIs to demonstrate that risk and validation functions are adequately resourced to provide genuine challenge to AI deployments, particularly where business lines have strong commercial incentives to deploy quickly. • Board AI education: MAS should expect Boards to receive periodic AI-risk training, similar to cyber-risk requirements in some jurisdictions. • Board-level AI or technology risk competence: In addition to general AI-risk education, MAS could recommend that the Board Risk Committee either (i) includes at least one member with substantive technology or AI risk experience, or (ii) formally retains an independent AI risk advisor who participates in relevant agenda items. This would help mitigate persistent information asymmetry and ensure that Board challenge to AI deployments is substantive rather than procedural. • Escalation pathways: When validation or risk functions identify concerns, is there a clear escalation path to the Board that bypasses business-line management? This should be explicit in the Guidelines. Illustrative edge case A CRO raises concerns about an AI model's fairness metrics, but the business line argues the model improves efficiency by 40%. The Board, lacking technical understanding, defers to commercial judgment. How does governance prevent this outcome? **Question 4:** Response Position: Support, with clarifications. I support the proposal for cross-functional oversight. However, structural design matters significantly.

		Key concerns • Risk-siloing concern: AI risk spans multiple existing risk domains (model, operational, conduct, technology). Without a cross-functional body, AI risks may fall between existing committees or be addressed inconsistently. I support MAS's approach of requiring cross-functional oversight where AI exposure is material. • Outcomes over structure: MAS should focus on outcomes rather than prescribing specific committee structures. FIs should demonstrate effective cross-functional coordination and clear accountability, whether through a new committee or enhanced existing structures. • Interaction with existing Model Risk Management: For banks, equivalents may already establish MRM governance. MAS should clarify how AI governance should integrate with (rather than duplicate) existing MRM structures. The allocation of responsibilities between an AI committee and existing model risk committees should be clearly defined. • Decision rights clarity: Can the AI committee veto a deployment? Impose conditions? Require remediation on a deadline? These decision rights should be specified. • Quorum and voting: For material decisions, should MAS expect documented voting or consensus, rather than "discussion and noting"? Assessing overall AI risk exposure MAS may wish to provide examples of criteria that would indicate material AI risk at the organizational level, such as: AI models materially influencing credit decisions, pricing, trading, or AML/CFT outcomes; high-volume customer interactions through AI-driven channels (e.g., chatbots, robo-advisors); or significant reliance on a small number of third-party foundation model or cloud providers for multiple critical use cases. FIs could be expected to assess their aggregate exposure against such criteria when deciding whether a dedicated cross-functional AI committee (or equivalent enhanced structure) is warranted. **Question 5:** Response Position: Support with modifications. I support comprehensive AI identification. However, practical challenges require attention. Key concerns • Shadow AI: Staff may use personal AI tools (ChatGPT, Copilot) without formal approval. Multi-channel identification approaches combining technical controls, procurement reviews, and attestation processes are necessary. • Vendor-embedded AI: Many existing vendor products have quietly incorporated AI features. FIs may not know they are using AI. Periodic vendor surveys and contractual notification requirements are essential.

		• API monetization and indirect AI exposure: If an FI provides APIs or data feeds to third parties who use AI on the FI's data, this creates indirect AI risk. The FI has limited visibility but potential reputational and legal exposure. MAS should acknowledge this as an emerging risk area. • Personal device usage: Staff increasingly use personal devices with AI assistants (Siri, Google Assistant, phone keyboards with AI). Where does identification stop? MAS should acknowledge the boundary problem. • Enforcement mechanisms: What consequences should attach to non-registration of material AI use? Without teeth, registration requirements become optional. Illustrative edge case A trader uses a personal phone with an AI assistant to dictate trading notes. The AI transcription sometimes "auto-corrects" numbers. Is this in scope? **Question 6:** Response Position: Support with modifications. I support comprehensive AI inventories. However, static inventories are insufficient for dynamic AI environments. I recommend MAS frame the inventory requirement as an AI Bill of Materials (AI BOM), analogous to software bills of materials (SBOMs) in cybersecurity. This framing emphasises that the inventory is not merely an administrative record but a living document that supports supply chain risk management, concentration risk analysis, and incident response. Key concerns • Dependencies and lineage: The inventory should capture dependencies between AI systems and upstream data sources. This enables concentration risk analysis and impact assessment when upstream components change. • Dynamism: GenAI models are frequently updated by providers. The inventory must capture version information and track changes, not merely record initial deployment. • Version pinning and reproducibility: For third-party LLM APIs that auto-update (OpenAI, Anthropic), FIs may lose the ability to reproduce results or understand behaviour changes. FIs should record whether they have the ability to pin model versions and, where they cannot, treat this as an additional risk factor. • End-of-life and sunset tracking: The inventory should flag models approaching end-of-support or sunset, especially for vendor-provided AI.

		• Contractual rights: The inventory should note whether the FI has contractual rights to audit, explain, or reproduce results. These vary significantly across vendors. • Concentration risk analysis: The inventory should support analysis of concentration risk across the AI portfolio (e.g., multiple critical systems dependent on a single foundation model provider). Proposed minimum inventory schema To operationalize the AI BOM concept, I recommend MAS specify minimum fields for AI inventories. At a minimum, each AI use case entry should include: • Use case name, accountable owner, and business purpose • Deployment context (internal or external; customer-facing or not) • Data types processed and sensitivity classification • Model and system components (model name, orchestration layer, retrieval store, tools) • Vendor details and contractual rights (audit, change notice, incident notice) • Versioning and update policy (pinned versus auto-updated) • Monitoring arrangements, key performance indicators, and incident history • Dependency links (upstream and downstream AI systems) **Question 7:** Response Position: Support, with additional dimensions proposed. I support the proposed dimensions of impact, complexity, and reliance. However, several additional factors warrant consideration. Additional risk dimensions To operationalise these dimensions, I recommend MAS endorse a simple scoring rubric. FIs could assess each AI use case against the following factors, with a “high” rating on any single dimension triggering enhanced controls, governance, and independent validation: • Impact (financial harm, customer harm, safety and soundness implications) • Autonomy (availability of human override and independence of review) • Scale (transaction volume and customer reach) • Reversibility (time to detect errors and ability to reverse or remediate) • Adversarial exposure (external access and incentive to attack or manipulate) • Dependency criticality (reliance on shared upstream components that could cascade) These can be treated as refinements or sub-factors within the existing dimensions of impact, complexity, and reliance, rather than as separate standalone axes, but they are important enough to merit explicit mention. • Reversibility and latency: How quickly can errors be detected and reversed? An AI system that makes irreversible decisions in real-time (e.g., automated trading) poses different risks than one where outputs are reviewed before action.

		• Volume and scale: The number of decisions affected. An AI system processing millions of micro-decisions daily may warrant higher scrutiny than one making a handful of significant decisions. • Data sensitivity: The nature and sensitivity of data processed by the AI system, including personal data, confidential client information, and market-sensitive data. • Adversarial exposure: Whether the AI system is exposed to adversarial actors who might attempt to manipulate its behaviour (e.g., customer-facing systems versus internal analytics). • Cascading and interdependency risk: If multiple AI systems depend on the same upstream AI component (e.g., a shared embedding model or a common LLM), a single failure can cascade. The risk materiality framework should account for this. FIs should assess not only the standalone risk of each AI use case but also its position in the FI's AI dependency graph. • Temporal concentration: Many AI models may be retrained or revalidated at similar intervals, creating "revalidation bunching" where validation resources are stretched. This is an operational risk that affects materiality assessment quality. • Gaming incentives: Without safeguards, FIs will have incentives to understate risk materiality. MAS should consider requiring independent validation of materiality assessments for high-impact domains. **Question 8:** Response Position: Support with modifications. I support comprehensive lifecycle controls. However, several GenAI-specific and emerging considerations warrant attention. GenAI and agent-specific recommendations • Red-teaming: For GenAI systems, structured adversarial testing (red-teaming) should be expected, including prompt injection testing, jailbreak attempts, and bias probing. • Sandboxing: AI agents with tool access should be tested in sandboxed environments before production deployment, with clear boundaries on permissible actions. • Kill-switches: High-autonomy AI systems should have robust mechanisms for rapid deactivation if behaviour exceeds risk tolerances. • Continuous monitoring for GenAI: Unlike traditional models with fixed inputs and outputs, GenAI behaviour can change with prompt variations. FIs should implement ongoing behavioural monitoring (e.g., sampling outputs for quality, safety, and consistency) rather than relying solely on periodic validation. Foundation model considerations • Tiered compliance: Where FIs fine-tune or deploy third-party foundation models, compliance obligations should be tiered based on the degree of customisation and the FI's ability to assess underlying model behaviour.

		Incident response • MAS notification: The Guidelines should clarify when AI-related incidents require notification to MAS. Quantitative thresholds (e.g., affecting more than X customers, more than $Y impact) would provide clarity. • Human-AI collaboration patterns: When multiple humans and AI systems collaborate in a workflow (AI drafts, human edits, AI refines, human finalises), accountability becomes diffuse. MAS should expect FIs to clearly designate accountability for final outputs in such workflows. • Model cards and documentation standards: MAS should consider referencing or developing documentation standards (similar to model cards) to ensure consistency across FIs. **Question 9:** Response Position: Additional considerations raised. I raise several additional considerations not fully addressed in the consultation. Workforce and capability • Key-person risk: AI expertise is concentrated among a small number of specialists. FIs should assess and mitigate key-person risk in AI functions, including succession planning and knowledge documentation. Systemic and sector-level risks • Systemic bias: If multiple FIs use similar AI models trained on similar data, sector-wide biases may emerge. MAS should consider macro-prudential monitoring of AI concentration and homogeneity. • Concentration risk at sector level: Over-reliance on a small number of foundation model providers creates systemic vulnerability. A major outage or security breach at a dominant provider could affect multiple FIs simultaneously. Customer-facing considerations • Customer transparency: Customers should be informed when AI materially influences decisions affecting them, consistent with emerging global expectations (e.g., EU AI Act). • Contestability: Customers should have meaningful ability to contest AI-driven decisions and receive human review where appropriate. Third-party and operational • Exit strategies and vendor lock-in: What happens when an FI wants to switch AI vendors? Are there adequate data portability and transition provisions? This should be explicit in third-party AI management. • Insurance and liability: Who bears liability when AI causes customer harm? How should FIs approach AI-related insurance? This is an emerging area worth flagging. Regulatory coordination

		• Cross-border coordination: The consultation is silent on how MAS will coordinate with other regulators (e.g., HKMA, FSC Korea, EBA) on AI governance. Given that many FIs operate regionally, some acknowledgment of this would be useful. • Interaction with existing MAS regulations: How do these Guidelines interact with existing MAS regulations on technology risk (TRM), outsourcing, business continuity, and fair dealing? FIs need clarity on overlaps and potential conflicts. **Question 10:** Response Position: Support, with practical recommendations. I support a 12-month transition period as reasonable for most FIs. However, practical implementation requires additional clarity. Recommendations • Milestone-based approach: Rather than a single 12-month deadline, MAS should consider a milestone-based approach. For example, AI inventory and gap assessment within 6 months, and full implementation status report at 12 months. Specifically, I recommend MAS require the following deliverables at each milestone: - Month 6: AI inventory complete, initial materiality ratings assigned, vendor disclosure survey completed - Month 9: Controls implemented for high-materiality AI systems, cross-functional governance committee (where required) stood up and operational - Month 12: Full compliance achieved, internal audit or independent assurance report completed (even if limited-scope) • Legacy system extensions: For AI systems deeply embedded in legacy infrastructure, case-by-case extensions may be necessary. MAS should establish a process for requesting extensions with appropriate justification. • New deployments during transition: New AI deployments during the transition period should comply with the Guidelines from deployment. However, clarity is needed for AI that is already in development at the time the Guidelines are issued. These “in-flight” projects need specific guidance. • Reporting requirements: What should FIs report to MAS during the transition period? I recommend that FIs submit an AI inventory and gap assessment to MAS within 6 months, and a full implementation status report at 12 months. • Enforcement posture: What is MAS’s expected enforcement posture during the transition? Will there be a “soft start” period where findings are observations rather than enforcement actions? Clarity would help FIs prioritise remediation efforts.
48	Lloyd's of London (Asia) Pte Ltd	**Question 1:** We support applying the Guidelines to all financial institutions on a proportionate, risk based basis. Proportionality should be anchored in factors such as AI autonomy, decision criticality, customer impact, scale of use, and degree of automation—with more stringent governance and continual testing for higher risk, business critical use cases (e.g.,

		underwriting/claims), and lighter touch controls for lower risk, assistive Generative AI under clear human oversight. The Annex is helpful, and additional examples of low risk internal use cases (e.g., document summarisation, productivity tools) would further promote consistent interpretation **Question 2:** We agree with a broad, technology agnostic scope that focuses on end to end deployments and outcomes/decision influence, not specific techniques. Inclusion of Generative AI and AI agents is appropriate, while purely deterministic, rules based systems should be excluded. We support framing around functionality and outcomes to keep the Guidelines future proof, and request clarity on how including "systems" adds value beyond governance applied to individual models. **Question 3:** We support MAS's allocation of responsibilities: the Board provides strategic oversight, approves the AI governance framework, and offers effective challenge without needing deep technical expertise; senior management is accountable for implementation and ongoing oversight. AI risk should be treated as cross cutting across existing risk types (model, data, conduct, operational, outsourcing, technology), rather than as a standalone silo, and applied proportionately. **Question 4:** We support cross functional oversight when overall AI risk exposure is material, with flexibility to leverage existing governance/risk committees so long as roles and accountability are clear. Overall exposure should aggregate individual use case materiality assessments while separately highlighting high risk cases; reviews should be dynamic. Adopting these practices early, even before AI use becomes material, can reduce duplication, promote consistency, and strengthen oversight as capabilities scale. **Question 5:** We support establishing clear definitions, criteria, and processes for consistently identifying AI across business and functional areas, supported by robust but proportionate systems that enable visibility and traceability. Guidance with concrete examples of expected definitions/criteria/processes would aid consistent implementation, and controls should avoid undue complexity for lower risk or assistive use cases. **Question 6:** We support maintaining an accurate, up to date AI inventory integrated within existing IT/solutions inventories to provide a holistic view and avoid duplication. Attributes should be decision useful and proportionate—e.g., purpose, risk rating, ownership, data sources, deployment status, and linkages to the broader technology landscape—to enable effective oversight and risk assessment. **Question 7:** We support focusing assessments on impact (with strong emphasis on customer outcomes) and likelihood, with complexity and reliance captured as key drivers. Institutions should have flexibility to treat complexity and reliance as standalone

		dimensions or as inputs to impact/likelihood; in all cases, periodic review is important to reflect evolving use and risk profiles. **Question 8:** We support comprehensive AI life cycle standards, processes, and controls, aligned to existing model/application life cycle frameworks where applicable. Controls should be applied proportionately to risk/materiality and level of autonomy, with a strong emphasis on ongoing monitoring, validation, and testing post deployment, while avoiding unnecessary duplication. **Question 9:** We request further guidance on how AI risk management interacts with outsourcing requirements, particularly where AI capabilities are embedded within vendor platforms, and clarification of supervisory expectations for internal productivity/assistive AI. Beyond these points, we have no additional comments. **Question 10:** We agree that a 12 month transition period is appropriate and sufficient given current levels of AI adoption and the limited number of active use cases.
49	Loh Wai Kwan	**Question 1:** A risk-based approach is more appropriate rather than proportioning to the size and nature of the FI's activities. A high-risk use of AI should warrant a tighter compliance to the guidelines. Examples of high-risk use of AI include operational processes that directly deal with customers' personal and financial data, executing monetary transactions, or decision support functions that may influence a high impact outcome etc; as compared to low-risk use of AI for example generating image for a new product launch, transcribing a meeting proceedings, etc. The application of the Guidelines should be focused on the High Risk, and perhaps including Medium Risk, use of AI. **Question 2:** The scope can be tier-ed into different risk categorises (eg. High, Medium, Low, or more) with different level of compliance to the guidelines for each category. High risk use of AI would include processes that directly deal with customers' personal and financial data, executing monetary transactions, or decision support functions that may influence a high impact outcome; Medium risk use could include market analysis, performance dashboarding, etc; Low risk could be image generation for product launch, transcribing meeting proceedings, etc. The scope should also be expanded to cover the agentic use of AI that can carry out executive orders on its own, and the risks arising from a networked deployment of AI agents/systems where AI agents/systems are connected and communicate with each other autonomously without human-in-the-loop. **Question 3:** The Board's oversight should be focused on the strategic deployment and policy of AI adoption and use, and not operational; operational matters are best managed by senior management to ensure operation within the boundaries and risk tolerances set by the Board. The Board risk oversight role should be subsumed under the mandate of an existing Board Risk Committee, or equivalent, who already has the mandate to oversee enterprise

		risks in the FI, rather than as a new Board Committee to minimise additional bureaucratic burden on the FI. **Question 4:** The proposed cross-functional committee should be a management-level committee and report to an existing enterprise risk management committee or equivalent. AI risk should be managed as part of the universe of enterprise risks faced by the FI, and not be silo'ed as a separate focus and to minimise additional bureaucratic burden on the FI. As such, the existing enterprise risk management framework of the FI should continue to be used for the assessment of overall AI risk exposure at the organisational level. **Question 5:** The definitions to facilitate the consistent identification of AI should be undertaken as an industry effort to standardise the understanding across the industry, perhaps led by MAS or IMDA. Once AI become ubiquitous, it will be very challenging to distinguish what is AI and what is not. **Question 6:** It may be futile to try to inventorise AI systems once the use of AI become ubiquitous in the FI. A risk-based approach may be better to only inventorise High Risk, and perhaps including Medium Risk also, AI usage. It may also not be a worthwhile endeavour to spend resources to establish and maintain an accurate and up-to-date inventory of all AI usage. **Question 7:** There are already various risk assessment methodology proposed by various government agencies on the assessment of risks for specific purpose, for example CSA guidelines on cybersecurity risk assessment for CII, and WSHA risk assessment for workplace safety and health etc. The similar but different methodologies of different agencies does not lend well for a FI to unify the management of enterprise risks when risks are assessed using different methodologies. It will be good if the guidelines are not too prescriptive to the point of being inflexible to align with existing methodologies of other government agencies and the internal enterprise risk management framework of the FI. **Question 8:** The life cycle controls should also include Replaceability and Removability. Once AI became entrenched in the architecture and infrastructure, it may be difficult to replace or remove it if such considerations were not incorporated at the onset. If an AI systems became obsolete, jailbreak'ed, starting to deviate from its original intended purposes, or even went rogue, there should be an avenue to either replace or remove it without major disruptions to the business activities it support. **Question 9:** The following should be considered and included: (1) Risks on the use of Agentic AI; (2) Risks arising from the deployment of a network of AI agents/systems with autonomous connectivity and communications with each other, without human intervention; (3) Risks of jailbreaking AI agents. Also, although the guideline is focused on managing the risks of AI, it may be good to incorporate a positive note to encourage the use of AI in the FI as this is the mega-trend towards the future.

		Question 10: At the current speed of AI advancement and innovation, 12 months is a long time.
50	Marsham Edge Pte Ltd	**Question 1:** We strongly support the principle of proportionate application as outlined in paragraph 1.5 and the Annex. A one-size-fits-all approach would disproportionately burden smaller financial institutions (FIs) and stifle innovation. We suggest the following enhancements to the guidance: - Clarify "Integrated Part of Business Processes": The examples in the Annex are helpful. We recommend further clarifying that low-impact, assistive AI tools (e.g., ChatGPT for drafting non-sensitive communications, grammar checkers, or research summarisation tools) used under human review should not automatically trigger the full suite of lifecycle controls. Their governance can be adequately covered by the "basic policies" tier. - Tiered Implementation Guide: Provide a simplified, tiered checklist or self-assessment tool for small FIs. This would help them map their AI use cases (Low/Medium/High Impact) to the corresponding expectations in Sections 2–5 of the Guidelines, making compliance more accessible and cost-effective. **Question 2:** The proposed scope is comprehensive. However, its practical application must consider a significant market barrier: Professional Liability (PL) Insurance. Insurance as a Constraint: For small firms and startups, obtaining affordable PL insurance for AI-related activities is exceptionally difficult. An overly broad or rigid interpretation of the scope could have the unintended consequence of: - Consolidation Risk: Driving AI development and deployment solely towards large, well-capitalised FIs that can absorb the risk, reducing market diversity and innovation. - Opacity Risk: Encouraging firms to use "black box" third-party AI services without proper diligence to avoid direct regulatory scrutiny, potentially increasing systemic risk. Recommendation: We urge MAS to engage with the insurance industry to encourage the development of accessible, fit-for-purpose AI liability products for smaller FIs. Furthermore, the Guidelines could explicitly recognise the use of certified or audited third-party AI platforms (aligned with MAS standards) as a risk-mitigating factor for smaller players. **Question 3:** We agree that the Board and senior management must own AI risk oversight. However, we note a critical capability gap. The Expertise Imperative: Board members cannot be expected to be experts in machine learning, reinforcement learning, or generative AI. Their role should be strategic oversight, not technical validation. Recommendation: The Guidelines should more strongly emphasise the Board's responsibility to ensure access to appropriate expertise. This could be fulfilled through: Appointing a dedicated Chief AI Officer or subject matter expert to the leadership team. Forming a technical advisory panel with external experts.

		Mandating that high-impact AI use cases are supported by documented reviews from qualified internal or external specialists. The Board's duty is to ensure this process is rigorous, not to understand the algorithmic nuances. **Question 4:** While we understand the intent, we are cautious about mandating a fixed committee structure, especially for smaller and medium-sized FIs. Agility Over Bureaucracy: The AI landscape evolves rapidly. A permanent committee can become slow and process-heavy. For many FIs, integrating AI risk into existing governance structures (e.g., Risk Committee, Technology Steering Group) may be more efficient and agile. Alternative Proposal: Instead of mandating a committee, the Guidelines could require FIs to demonstrate "effective and documented oversight" for material AI risks. This could be achieved through: - A dynamic "AI Risk Review Forum" convened as needed for specific high-impact initiatives. - Clear documentation of the oversight model chosen (committee, integrated, or hybrid). A mandatory, maintained register of high-impact AI models with a documented annual review and sign-off by the Board or its delegate, focusing on the model's purpose, potential impact, and key risk indicators. **Question 5:** We support the need for identification and an inventory but urge pragmatism in its scope. Pragmatic Inventory: Requiring an inventory of every AI use, including every instance of an employee using an AI-powered spelling checker, is administratively burdensome and adds little risk management value. Risk-Based Triage: The process should focus on systematic, integrated, or high-impact AI. We propose a two-tier inventory: - Tier 1 (Monitored): Low-risk, assistive tools governed by company-wide acceptable use policies (e.g., approved generative AI tools with data guards). - Tier 2 (Inventoried & Managed): AI models or systems that influence customer outcomes, automate significant processes, or handle sensitive data. These require full documentation in the register (purpose, owner, risk rating, validation status). **Question 6:** See Response to question 5. Not all AI usage needs to be managed in our view. **Question 7:** The proposed dimensions of Impact, Complexity, and Reliance are sound. We suggest one addition: - Velocity & Adaptability: Consider the speed at which the AI model learns or changes (e.g., online learning models, frequently re-trained agents) and the adaptability of its outputs. A high-velocity, highly adaptive model in a customer-facing role may pose unique conduct and stability risks not fully captured by "complexity" alone. **Question 8:**

		The listed controls are comprehensive. The key to their success is the proportionate application reiterated throughout. We suggest the final Guidelines include illustrative examples pairing specific use cases (e.g., "Chatbot for general FAQs" vs. "AI for credit scoring") with the expected minimum control sets. This would provide invaluable clarity, especially for smaller FIs building their frameworks. **Question 9:** An area not explicitly covered is the management of "Shadow AI", or the unsanctioned use of AI tools by employees. Recommendation: The Guidelines could encourage FIs to: - Educate staff on risks and approved tools. - Provide secure, sanctioned alternatives to popular public tools. Implement technical controls (where feasible) to detect and manage unauthorised usage, framing it as an operational risk issue. **Question 10:** A 12-month transition period is reasonable for FIs to map their existing landscape and establish foundational policies. We recommend that MAS clarify that during this period, the expectation is for FIs to develop and approve their implementation plans, with a focus on addressing high-impact use cases first. We are happy to provide further clarification on any of the points raised and welcome continued dialogue on this critical topic. We work with family offices, governmental agencies, energy & construction companies to advise them on their AI journey.
51	Masato Tsukada	**Question 1:** **Question 2:** **Question 3:** **Question 4:** **Question 5:** **Question 6:** **Question 7:** Additional risk dimensions to be included in AI risk materiality assessments The proposed risk dimensions of impact, complexity, and reliance are appropriate and effective for assessing many classes of AI-related risks. They provide a strong foundation for evaluating the potential consequences, technical sophistication, and degree of dependency associated with AI systems. However, with the increasing deployment of agentic AI systems and continuously updating or learning models, an additional category of risk has become more prominent: the risk of gradual deviation from originally intended objectives over time.

		To address this, it is recommended that a fourth, complementary risk dimension be considered alongside the existing three: Alignment Risk (Cognitive Drift). Cognitive Drift refers to the phenomenon whereby an AI system progressively deviates from its originally intended objectives, decision criteria, or normative constraints as a result of ongoing interaction with its environment, data updates, or continuous learning processes. This type of risk is inherently time-dependent and is difficult to detect through single-point performance testing or static audits alone. As supplementary signals for AI risk materiality assessments, the following observable indicators may be considered: Temporal changes in output distributions under identical or comparable input conditions Emerging signs of misalignment from initial objectives or policy constraints, including mission creep Variations in the frequency or pattern of human intervention and escalation events The specific measurement approaches and threshold settings for these indicators should be designed flexibly, in accordance with each financial institution's risk appetite, business context, and operational characteristics. For example, a credit risk assessment AI may initially be trained to produce conservative default probability estimates. Over time, through continuous learning during a prolonged economic expansion, the model may gradually shift towards increasingly optimistic estimations, eventually deviating from established risk management standards. This represents a typical manifestation of Cognitive Drift. Incorporating Alignment Risk as a complementary risk dimension would enhance the ability of financial institutions to identify and manage time-dependent risks specific to advanced and agentic AI systems, while remaining consistent with international approaches to AI risk management. **Question 8:** Standards, processes and controls across the AI lifecycle AI risk management standards, processes, and controls should be applied consistently across the entire AI lifecycle, including design, development, deployment, operation, and audit. These controls should be applied in a proportionate manner, taking into account the specific AI model, system, and use case, as well as the associated risk materiality. In the context of agentic AI systems and models that adapt or update over time, lifecycle governance should address not only decision-making behaviour, but also the structured deferral of decisions when uncertainty or misalignment risks exceed predefined thresholds.

		One important lifecycle control is the standardisation of decision deferral mechanisms, referred to as a Silence Event. A Silence Event occurs when an AI system refrains from making an autonomous decision and escalates the case to human oversight after predefined safety thresholds are breached. Such thresholds may include confidence measures, prediction errors, or alignment evaluation signals. This mechanism supports safe operation by preventing potentially erroneous or harmful decisions under conditions of elevated uncertainty. In addition to static performance testing, AI lifecycle controls should incorporate dynamic and continuous risk evaluation. Relevant areas for assessment include: Data drift monitoring, to detect changes in input distributions over time Concept drift detection, to identify changes in the relationship between inputs and outputs Alignment evaluation, to enable early detection of deviations from intended objectives or policy constraints These dynamic monitoring and deferral mechanisms align with the EU AI Act's requirements for human oversight and continuous monitoring of high-risk AI systems, as well as the NIST AI Risk Management Framework's emphasis on ongoing measurement and management of AI risks throughout the system lifecycle. During pilot or phased deployment, enhanced controls are particularly important. These may include clearly defined limitations on duration and scope of use, predefined success criteria and termination conditions, and explicit transition requirements for progression to full-scale production. Decision deferral behaviour during these phases should be formally evaluated within existing fiduciary responsibility frameworks. Operational delays resulting from appropriately designed and documented decision deferral should be recognised as legitimate safety measures intended to prevent erroneous or harmful outcomes. Such deferral may be positively evaluated in supervisory contexts when: (a) the rationale for deferral, including confidence, error, or alignment indicators, is clearly recorded; (b) the reasonableness of threshold settings is documented in advance; and (c) post-event review and improvement processes are appropriately conducted. For example, when a credit approval AI detects that an applicant's attributes fall outside the distribution of its training data, escalating the case to a human reviewer rather than issuing an automatic approval or rejection represents a proportionate application of lifecycle controls through a Silence Event. The effective implementation of these lifecycle controls, particularly the design and continuous adjustment of Silence Event thresholds and escalation pathways, may benefit from dedicated governance expertise, as discussed in the response to Question 9.

		Question 9: AI risk management aspects not sufficiently covered by the guidelines Building on the lifecycle controls outlined in Question 8, particularly the use of Silence Events and dynamic risk evaluation mechanisms, this response addresses the human governance function required to design, implement, and supervise these controls effectively. In environments involving highly autonomous or multi-agent AI systems, traditional Human-in-the-Loop approaches that rely on human approval of individual decisions face practical scalability limitations. As AI systems operate at greater speed, scale, and complexity, human oversight must evolve from direct decision approval to meta-level governance. This specialised governance function may be formalised as a distinct role, such as an Autonomous Systems Governance Architect, or integrated into existing governance structures depending on institutional scale, complexity, and risk profile. The purpose of this function is to address technical and operational characteristics specific to advanced and agentic AI systems, rather than to replace existing risk management or compliance roles. In smaller institutions, these responsibilities may be distributed among existing functions such as risk management, model validation, and technology governance. In larger institutions deploying multiple autonomous systems, a dedicated function may be warranted to ensure consistent and effective oversight. Key responsibilities of this governance function may include: Designing and validating interaction rules and alignment across multiple AI agents or systems Defining and adjusting decision deferral and escalation conditions, including Silence Event triggers and threshold configurations Designing escalation pathways and ensuring that relevant contextual information is provided to human supervisors Continuously monitoring indicators of Cognitive Drift and reflecting these observations in AI risk materiality assessments This governance function supports the implementation of meaningful human oversight as required under the EU AI Act and contributes to the development of organisational competencies necessary for trustworthy AI deployment in accordance with international best practices. For example, in a large financial institution deploying multiple AI agents across customer service, risk management, and fraud detection functions, this function would oversee the

		integrated design and supervision of action boundaries, interaction rules, and escalation conditions across these systems. Clarifying this governance function would strengthen the practical implementation of AI risk management, enhance proportionality across institutions of different sizes, and improve the effectiveness of existing AI governance frameworks. **Question 10:**
52	MindXO Consulting FZCO	**Question 1:** We support the proposed proportionate application of the Guidelines. The principle that governance intensity should scale with risk materiality is both pragmatic and aligned with international best practice. We offer the following observations to strengthen implementation: 1. Criteria for Assessing AI Adoption Level The Guidelines establish that "All FIs should minimally institute basic policies for the use of AI commensurate with the FI's level of AI adoption" (para 1.5) while FIs with AI integrated into business processes require fuller governance. However, the criteria for assessing an FI's level of AI adoption are not specified. We suggest that the AI Risk Management Operationalisation Handbook provide guidance on how FIs may self-assess their AI adoption level. Such guidance would help FIs determine when their AI adoption has progressed beyond "minimal" and triggers the expectation for enhanced governance structures. This would also provide MAS with a more consistent basis for understanding AI adoption patterns across the financial sector over time. 2. Sandbox Exception for Experimental AI Use The Guidelines apply to FI whether or not AI is used, with more stringent expectation if AI is used as an integrated part of business processes. However, there is a third category that does not appear to be explicitly addressed: controlled experimental use of AI in sandbox or pilot environments. We suggest the Guidelines to clarify that FIs may include provisions for sandbox or pilot AI use as a special case within their basic policies. Such provisions would govern AI systems that: • Are deployed in a controlled test environment, segregated from production systems • Are not used for actual business decisions affecting customers or operations • Are limited to a defined pilot population or synthetic data • Are subject to a defined time limit for experimentation • Are not connected to critical business processes or customer-facing channels This clarification would provide FIs confidence that carving out experimental use within their basic policies is an acceptable approach, while ensuring appropriate safeguards remain in place (e.g., data protection, security controls, defined exit criteria). The transition from pilot/sandbox to production deployment should be a defined governance checkpoint, at which point the Guidelines would apply based on the AI system's risk materiality.

		3. Guidance for Smaller FIs While proportionality inherently accommodates smaller FIs, explicit guidance would assist implementation. FIs with minimal AI use might satisfy requirements through: a single AI policy covering all use cases, simplified intake and approval processes, combined roles (e.g., one individual serving as both Model Owner and Risk Owner for low-risk AI), and reliance on vendor documentation for third-party AI with appropriate due diligence. 4. Supporting Implementation Through Standard Templates To facilitate consistent implementation across the industry, we recommend MAS consider providing standard templates for key artefacts at a later stage, for example through the AI Risk Management Operationalisation Handbook. This could include templates for AI policies, intake forms, risk assessment frameworks, and evidence documentation. Standard templates would reduce implementation burden, particularly for smaller FIs, while promoting consistency in how the Guidelines are applied across the financial sector. **Question 2:** No comment. **Question 3:** We support the proposed Board and senior management responsibilities. We offer the following observation for MAS consideration: Additional Guidance on Role Definitions The Guidelines appropriately require that the AI inventory capture "key roles and responsibilities (e.g., owners, developers)" (paragraph 3.5) and that clear roles be assigned for AI risk materiality assessment, including a control function to ensure consistent application (paragraph 3.11). While these requirements establish sound foundations, FIs would benefit from additional guidance on the types of roles to be defined and the distinction between them - particularly between risk management and governance functions. For example, clarity on the respective responsibilities of an AI Risk Owner (accountable for risk assessment and monitoring) versus an AI Governance Owner (accountable for policy compliance and approval processes) would support consistent implementation across the industry. The Guidelines could also clarify how role assignments and approval mechanisms should scale with risk materiality, both at the individual AI system level and at the organisational level. For FIs with lower risk materiality, a single individual may appropriately cover multiple roles and approval processes may be streamlined. As risk materiality increases, roles should be separated and approval mechanisms strengthened to ensure adequate oversight and avoid conflicts of interest. Such guidance could be provided through the AI Risk Management Operationalisation Handbook, allowing FIs flexibility while promoting consistency in how roles are defined and scaled across the financial sector. **Question 4:**

		We support the expectation for a dedicated cross-functional committee where organisational AI risk exposure is material. This is one of the most important governance mechanisms for ensuring coherent oversight across what are often fragmented AI initiatives. We offer observations on both the committee structure and the materiality assessment methodology. Part A: Cross-Functional Committee Structure We agree with the requirement for a cross-functional committee but suggests the Guidelines clarify that the committee should operate at two levels: Strategic Oversight: Responsible for AI strategy, risk appetite, policy approval, and escalation of material matters to Board and senior management. This level ensures alignment between AI initiatives and organisational objectives. Operational Implementation (can be tackled in AI Risk Management Operationalisation Handbook): Responsible for reviewing and approving individual AI use cases, monitoring ongoing compliance, and coordinating across business and control functions. This level ensures governance is applied consistently at the system level. Clear roles should be defined at each level to avoid duplication and ensure accountability. The strategic and operational functions may be discharged by the same committee for FIs with lower AI risk exposure, but should be clearly distinguished as risk exposure increases. Part B: Assessing Organisational-Level AI Risk Materiality Notwithstanding our comments on paragraph 3.10 in our response to Question 7, we offer the following observations on assessing organisational-level AI risk exposure. We agree with the guidance in footnote 12 of paragraph 2.4 that the existence of one or more AI use cases exposing the FI to significant risks should trigger assessment of whether a dedicated cross-functional committee is necessary. In our view, the presence of at least one high-risk AI system deployed in critical business lines or regulated activities should generally warrant establishment of such a committee. Beyond individual high-risk systems, FIs could consider additional factors when assessing organisational-level exposure: Aggregation Across Use Cases: The number and distribution of AI use cases across risk materiality tiers provides a starting indicator. An FI with numerous moderate-tier AI applications may have material aggregate exposure even without any single high-tier application. We suggest organisational materiality thresholds such as: any critical-tier AI in production, at least [X] high-tier AI applications, or [Y] or more moderate-tier AI applications in aggregate. With thresholds self-determined by the FI's size and risk profile. Concentration Risk (extension of the concentration risk described in para 4.11 of the Guidelines): Multiple AI applications in the same business line or serving the same customer segment may create correlated failure risk. Multiple AI applications relying on a single vendor or underlying model may create single points of failure. Customer Exposure Metrics: The proportion of customers or transactions affected by AI-driven decisions indicates organisational exposure. Relevant metrics may include

		percentage of credit decisions influenced by AI, percentage of customer interactions handled by AI, and value of assets managed with AI-assisted decisions. Dependency and Reversibility: Consider how deeply AI is embedded in operations. Could the FI continue critical operations if AI systems were unavailable? How quickly could AI-driven decisions be reversed if systematic errors were discovered? Can a kill-switch be implemented to halt AI operations if required, and how rapidly? Velocity of AI Adoption: Rapid AI adoption may itself indicate material exposure, as governance maturity may lag deployment. FIs experiencing significant growth in AI use cases should assess whether governance capabilities have scaled proportionately. **Question 5:** Shadow AI Discovery for FIs with Minimal AI Adoption The Guidelines appropriately establish robust AI identification requirements for FIs with AI integrated into business processes (Section 3- AI Identification), including the establishment of "systems, policies and procedures to ensure the consistent identification of AI usage across all relevant business and functional areas. " We understands that this addresses "shadow AI" risk for FIs at that adoption level only (e.g. with AI integrated into business processes). However, for FIs with minimal or no known AI adoption - those applying only basic policies (Annex, paragraph 5)- there is no equivalent requirement to proactively discover unreported AI usage. The basic policies (Annex, paragraph 5) require periodic compliance checks, but these assume the FI already has visibility over where AI is being used. This creates a gap: an FI that believes it does not have AI integrated unto business processes may have shadow AI that, if discovered, would reclassify the FI as having AI integrated into business processes - triggering fuller governance requirements. We suggest the basic policies to include an explicit expectation for periodic discovery of AI usage, proportionate to the FI's size and complexity. This needs not be as rigorous as the Section 3 identification systems, but should ensure FIs do not inadvertently remain in the "basic policies only" tier when their actual AI footprint warrants enhanced governance. **Question 6:** No comment. **Question 7:** We support the proposed three-dimensional framework of Impact, Complexity, and Reliance for assessing AI risk materiality. This framework provides a sound basis for proportionate governance. We offer observations on each dimension. Dimension 1: Impact The Guidelines identify relevant impact factors including financial, operational, regulatory, reputational, fairness, ethical breach and consumer protection. To support consistent and

		comprehensive assessment, FIs would benefit from a structured taxonomy of impact domains. While the Guidelines appropriately remain high-level, further elaboration of these factors in the Operationalisation Handbook for instance would support consistent application. We propose the following: 1. Strategic & Alignment: The risk that AI systems are framed or deployed in ways that do not align with strategic objectives or intended purpose. 2. Commercial & Customer: Risk that AI systems negatively affect commercial performance, customer value, contractual obligations, or market competitiveness. 3. Operational & Business: The risk that AI deployments impair business operations and continuity, workflows, productivity, or service performance. 4. Regulatory & Legal: The risk that AI systems violate legal or regulatory requirements due to insufficient governance. 5. Health, Safety & Environment (HSE): Risk that AI systems cause or contribute to physical harm, safety incidents, or environmental damage. 6. Reputational & Brand: The risk that AI system outputs, behaviour, or failures undermine stakeholder trust, damage the organisation's public image or brand equity, or erode confidence in its products, services, or leadership. 7. Financial: The risk of financial loss due to inefficient AI spending, penalties, operational failures, or poor vendor decisions. 8. Confidentiality & IP: Risk that AI systems leak, misuse, or compromise confidential information or intellectual property. 9. Technical Performance & Reliability: The risk that AI systems fail to perform reliably or consistently due to drift, insufficient validation, or inadequate robustness. 10. Security & Resilience: The risk that AI systems are compromised or degraded due to vulnerabilities, attacks, or lack of resilience. 11. Data & Privacy: The risk that AI systems cause privacy violations, data misuse, or data quality failures. 12. Algorithmic Ethics, Fairness & Society: The risk that AI systems produce harmful, unfair, biased, discriminatory, or socially adverse outcomes due to issues in data, modelling, or algorithmic behaviour. 13. Transparency & Traceability: The risk that AI systems cannot be understood, interpreted, audited, or meaningfully explained by stakeholders, leading to misuse, wrong decisions, compliance failures, or erosion of trust. A structured taxonomy ensures FIs assess impact comprehensively rather than focusing only on the most obvious risks. It also supports aggregation and comparison across AI systems within the organisation. Dimension 2: Complexity The proposed complexity factors (nature of AI technology, novelty of application, data complexity) appropriately capture technical risk drivers. We suggest FIs additionally consider: • Model Opacity: The degree to which model behaviour can be understood and explained. Black-box models warrant higher complexity classification than interpretable models. • Integration Complexity: How deeply the AI is integrated with other systems. Standalone AI presents different risks than AI embedded in critical transaction processing chains.

		• Novelty Within the FI: Even well-established AI techniques may be novel to a particular FI. Organisational inexperience with a technology type elevates complexity risk regardless of market maturity. • Change Velocity: The frequency and nature of expected changes to the AI system, including retraining cycles, vendor-initiated updates, and data source changes. High-change-velocity AI requires more intensive ongoing governance. Dimension 3: Reliance The proposed reliance factors (level of autonomy, degree of human oversight, availability of alternatives) are well-chosen. We suggest FIs additionally consider: • Override Practicality: Nominal human oversight is ineffective if overriding AI recommendations is practically difficult due to time pressure, information asymmetry, or organisational culture. Assessment should consider whether override is genuinely feasible. • Kill-Switch Capability: Whether the AI system can be rapidly halted if critical failures are detected, and how quickly operations can revert to non-AI alternatives. • Expertise Degradation: Over time, reliance on AI may erode human expertise needed to operate without AI. This creates path dependency that increases effective reliance beyond what current assessment might suggest. AI Risk Management Operationalisation Handbook To ensure consistent application, we suggest that FIs develop scoring methodologies for each dimension, with explicit criteria for each score level. Aggregation across dimensions should be defined (e.g., highest dimension determines tier, or weighted combination). This prevents inconsistent materiality assessments across similar AI applications and enables portfolio-level risk aggregation. Should MAS find it useful, the observations above could be further developed into practical examples or implementation artefacts in support of the forthcoming AI Risk Management Operationalisation Handbook. **Question 8:** No comment. **Question 9:** We commend MAS for developing comprehensive Guidelines that appropriately balance principles-based expectations with practical implementation flexibility. We offer the following additional observation on aspects that may warrant consideration in the Guidelines or the forthcoming AI Risk Management Operationalisation Handbook. 1. AI Incident Classification and Learning We suggest the Guidelines or Handbook address AI-specific incident classification. Traditional incident management may not adequately capture AI failure modes such as gradual drift, emergent bias, or adversarial exploitation. A classification framework distinguishing AI-specific incident types would support more effective root cause analysis and industry learning. Over time, such classification could also support anonymised industry-level learning and supervisory dialogue on emerging AI risk patterns.

		Question 10: No comment.
53	MUFG Bank, Ltd.	**Question 1:** Given the diversity in AI adoption across FIs, it is important that requirements scale with the size, complexity, and risk profile of each FI. We suggest providing clear examples or case studies to assist FIs in applying the proportionality guidance effectively, enabling them to focus resources on areas of highest risk and impact. **Question 2:** We agree with the proposed scope, encompassing machine learning, deep learning, Generative AI, and AI agents. It is important to distinguish between AI systems that materially affect business or customer outcomes, and those used for assistive, non-decision-making purposes. This will ensure risk management efforts are focused where most needed. **Question 3:** With reference to Section 5 of the Annex of the Proposed Guidelines, we seek MAS' clarifications on the need to designate a specific Senior Management member for AI oversight, as opposed to adopting a general governance approach such as oversight via an existing management committee. This requirement might be too onerous for a basic policy, especially when the level of adoption of AI is low. **Question 4:** NIL **Question 5:** We seek MAS' clarification and guidance on the requirement to have robust systems to support the AI Identification process, especially for FIs with non-extensive count of AI implementation, or largely lower risk AI implementation. We request that MAS allow FIs to scale according to the volume and risk for AI Identification and broader Risk Management. Identification of all AIs may no longer be feasible or reasonable. The proliferation of AI across all services have resulted in environments where there are constant exposure to AI whether known or unknown – such as shadow AI or AI functions added to SaaS tools without prior notice etc. We request MAS to provide more guidance on managing such AIs. **Question 6:** We seek MAS' clarification and guidance in detailing the compulsory or minimum fields (e.g., purpose, risk rating, data sources, validation status) to be tracked in the AI inventory for different types of AI (AI LLM dev, AI Tool purchase, AI used by third party), depending on their necessity and risk. **Question 7:**

		We seek MAS' clarification and guidance (e.g. via case studies, assessment templates) in detailing and defining what are the matrix/levels of complexity and reliance to assess AI risks materiality. We also suggest MAS to consider who trained the AI model, as an AI model trained by external AI vendor will limit the information we can obtain in terms of its AI model development, testing, deployment. **Question 8:** There may be limited information provided to FIs by third party AI providers due to practical constraints (e.g. inadequate disclosure of information), hence it may not be possible to implement necessary mitigants nor limit the usage of AI by our 3rd parties. We suggest that MAS consider including workarounds for such scenarios or having independent auditors to assess the third party AI risk posture (e.g. OSPAR, SOC2). While Section 4.23(b) stipulates that resolution methods should depend on the severity and nature of an issue, we seek clarification on whether MAS intends to specify severity thresholds or classification frameworks, together with corresponding control or response requirements. Under Section 4.12 on 'Selection', it is stated that the developers are required to justify and document their selection process. As the responsibilities stated may not be specific to developers only, we suggest broadening the requirement to avoid limiting it to developers only. **Question 9:** Section 5.2 on 'AI Capability & Capacity' states the need for regular reviews on the relevant personnel involved in developing, deploying and maintaining an AI use case. The current wording suggests that a formal review (and documentation) is required regularly, which may be onerous. We suggest adopting a less prescriptive approach and provide flexibility to FIs to decide on the appropriate manner to keep their AI risk management capabilities up to speed. We also seek MAS' support to address emerging risks such as adversarial attacks, model drift, and ethical considerations in Generative AI. Guidance on incident management, kill switches, and redress mechanisms for customers affected by AI decisions would be valuable. **Question 10:** During the transition period, we suggest that MAS consider providing industry support measures such as FAQs, best practice publications, or workshops to facilitate smooth implementation of the Guidelines.
54	Omnitrade Technologies Pte Ltd	**Question 1:** Omnitrade is a Singapore-based technology provider that operates institutional-grade trading and operations infrastructure used by buy side firms globally. Our platform is embedded in market-critical workflows, including order and execution management, broker connectivity, portfolio oversight, market data ingestion, and regulatory reporting. As a small but growing firm, we have undergone extensive due diligence and security/operational reviews as part of the vendor approval process of many regulated financial institutions. This gives us direct exposure to how regulatory expectations are

		interpreted in practice, and how those expectations resultantly shape competition and innovation. We support MAS' proposal to apply the Guidelines proportionately, particularly the distinction between assistive AI and AI integrated into business processes. Proportionality should be driven less by the size of the institution and instead by how deeply AI is embedded in day-to-day workflows and decision-making. While senior management may endorse AI initiatives or highlight pilot programmes when presenting at conferences, front-office adoption is often cautious. In our experience, this mirrors early resistance seen during the shift to electronic execution and algorithmic trading. That caution is typically driven by concrete issues: limited transparency into model behaviour, perceived loss of control and uncertainty around accountability. Infrastructure platforms can help with this by applying lightweight controls where AI is used in an assistive capacity and progressively stronger governance where AI meaningfully influences outcomes. This outcome-based segmentation reflects practical workflows and supports responsible experimentation without introducing unnecessary friction. In many institutions, there is a large disconnect between stated AI strategies and readiness to deploy AI within live production environments. Technology providers therefore play an important role in translating proportional supervisory intent into practical, risk-aligned controls and at the same time, creating clean and auditable data trails that improve visibility into how AI is used. **Question 2:** **Question 3:** **Question 4:** We support MAS' proposal for cross-functional oversight where overall AI risk exposure is material. However, cross-functional committees may not be practical or proportionate for smaller financial institutions or those with limited AI usage. AI-related risks frequently cut across organisational boundaries; front-office teams may rely on execution algorithms without visibility into data dependencies, meanwhile risk teams may consume AI-driven analytics without awareness of model changes. Additionally, compliance teams may find it difficult to identify AI usage when it is embedded within third-party platforms that don't notify such changes. In practice, effective cross-functional committees require sustained engagement, senior representation, and clear authority across business lines. For smaller institutions, this level of overhead may outweigh the risks posed by early-stage or low-impact AI use. Effective oversight depends on both the visibility into where AI is used and the internal capability to assess the associated risks. While MAS expects boards to provide input and well-informed oversight, many institutions still lack this capability. As a result, boards may either approve AI deployments without a full appreciation of their risk profile or decline valuable use cases due to unfamiliarity. Overall AI risk exposure could be assessed by looking across all individual AI use cases rather than treating it as a binary threshold by aggregating use-cases across the

		organisation. This could be carried out by conducting assessments of each use case, evaluating presence in critical workflows, the degree of reliance placed on AI outputs, and whether such results influence customer/market outcomes. Risk can become material not only through a single high-impact deployment, but also through the cumulative effect of many low-risk use cases throughout an FI. For smaller institutions with limited AI deployment, this assessment could be owned by a designated senior executive and reviewed periodically, similar to existing accountability frameworks such as the appointment of a Data Protection Officer under the PDPA. This simplifies matters rather than requiring an entire committee for those with limited resources. Nevertheless, this still relies on clear visibility of AI usage throughout the organisation which may not always be possible. Infrastructure platforms can support more comprehensive risk assessments through their knowledge of the operating practices of firms across the industry as well as in-house technological expertise. MAS may wish to recognise the current variation in institutional capability, particularly among smaller firms that cannot readily justify dedicated AI governance roles. **Question 5:** We support MAS' emphasis on consistent AI identification. In our experience, this is a foundational requirement and one of the most underestimated aspects of AI risk management. Implementation, however, is not so straightforward. As AI functionality becomes embedded within wider systems, its presence is not always visible at first glance to end users. Execution platforms, data pipelines, and workflow tools may incorporate adaptive logic without explicit labelling. Business users may not know whether AI is involved, and technology teams may have limited insight into third-party implementations. This increases the likelihood of incomplete inventories. **Question 6:** We support MAS' emphasis on accurate AI inventories and strongly believe trading and operations platforms can help to address this. Platforms routinely maintain end-to-end records of order flow, broker connectivity, and regulatory reporting in common standards such as FIX. They can also tag AI use cases and link inventory records directly to execution and post-trade processes in a way that improves traceability. MAS may wish to consider how technology inventories are currently maintained across the industry. Many institutions still rely on spreadsheets or informal knowledge held by critical individuals. Few have comprehensive system inventories today. Expectations for "robust systems" therefore represent a meaningful step change, particularly for smaller institutions. **Question 7:** **Question 8:** We support MAS' expectation that AI life cycle controls should scale with risk. In trading-related contexts, the emphasis is usually on a small number of practical controls. This includes keeping assistive tools separate from decision-making systems, managing changes carefully, being able to review AI-driven activity when needed, and having workable fallback arrangements for market-critical functions.

		Infrastructure platforms can apply these safeguards consistently across multiple use cases, reducing fragmented implementation. However, some proposed life cycle controls assume levels of AI maturity and internal expertise that are not yet common across the industry. Controls such as independent pre-deployment validation, formal fairness assessments, or detailed explainability analysis may be appropriate for high-impact, customer-facing AI deployed at scale. Applying the same requirements to low-risk internal tools is unlikely to be proportionate and may discourage experimentation without materially improving outcomes. Illustrative examples of proportionate life cycle controls across different use cases would help institutions calibrate their approach. Without such guidance, there is a risk that low-risk deployments are subject to excessive controls, while higher-risk use cases do not receive sufficient focus. **Question 9:** Third-Party AI Risks: Accountability and Opacity We welcome MAS' focus on third-party AI risks, including opacity, dependency, and concentration. This is an area where infrastructure platforms can provide value, but where unintended consequences should be carefully considered. Financial institutions are often accountable for outcomes without having meaningful leverage over large AI providers. Dominant vendors may offer limited disclosure, particularly to smaller customers or downstream technology partners. Third-Party AI Risks: Vendor Capability and Exclusion Risk At the same time, smaller AI vendors may lack the resources to meet enterprise-grade documentation or validation expectations. Where supervisory expectations implicitly assume uniform vendor capability, smaller providers may be excluded. Infrastructure platforms can mitigate some of these risks by abstracting underlying AI providers, enforcing usage constraints, and providing consistent monitoring and logging across integrations. This allows institutions to detect unexpected behaviour even where full model transparency is not available. Proportional Application of Controls Across Use Cases We encourage MAS to clarify that third-party AI risk management should reflect both the materiality of the use case and the nature of the vendor relationship. A customer-facing chatbot built on a large language model warrants extensive due diligence. By contrast, infrastructure software that applies adaptive optimisation for internal trade routing may be better assessed through controls focused on operational resilience, change management, and cybersecurity rather than detailed model validation. Competition, Concentration and Industry Coordination Compliance expectations that assume extensive documentation, formal validation processes, and enterprise-grade support infrastructure may disadvantage smaller technology providers. Over time, this could reduce competition and increase concentration risk. Inconsistent AI identification and inventory practices across vendors make enterprise-wide visibility difficult. Industry collaboration on shared taxonomies or data models would materially improve risk management outcomes. We as an organisation have paid close attention to the work of the AI Working Group (AIWG) of the FIX Trading Community, in respect to its recent work on AI identifiers for trading.

		Omnitrade supports MAS' proposed Guidelines as an important step toward responsible AI adoption in the global financial sector. We encourage continued attention to how proportionality is interpreted in practice, particularly in relation to smaller institutions and technology providers like ourselves. Infrastructure platforms play a critical role in translating supervisory expectations into operational controls. Ongoing engagement between regulators, financial institutions, and vendors is imperative to ensure that AI risk management strengthens resilience without limiting industry participation or innovation. We would welcome the opportunity to contribute further through industry working groups, sandboxes, or technical discussions. **Question 10:** We support the proposed 12 month transition period which should provide sufficient time for most institutions to assess and implement the Guidelines. Although it is worth noting that smaller firms operating on legacy infrastructure may require additional guidance during this period.
55	Oversea-Chinese Banking Corporation Limited	**Question 1:** We welcome the use of the principle of proportionality in the application of the AIRG. With reference to the Annex in the Consultation Paper, we would propose that this principle be applied at the use case-level, rather than at the institution-level, such as in the deployment of Gen AI. We would like to seek MAS's guidance on whether the applicability of the Guidelines includes subsidiaries in other countries whose parent entity is in Singapore. From the Annex, we seek MAS's consideration to provide clearer guidance. Some of the following points could result in differing interpretations: (1) Proportionality is unclear given the language used, such as "significantly," "much larger," versus "delay claims," and "reduce the ability." Should there be clearer guidance to ensure consistency in gauging what is considered integral? (2) There was an example in section 3(f) stating that ad-hoc use is not regarded as using AI; however, if similar AI usage occurs periodically, such as for initial reviews of claims documents, would that then be considered using AI as an integral part of business processes? **Question 2:** We agree with the proposed definition and scope of AI use cases, systems and models for the applications of the Guidelines. However, we are cognizant of the challenge to identify AI components, especially those embedded within vendor products/service. To what extent would this Guideline subject third party service providers to disclose their AI usage in their application? **Question 3:** We largely agree with the proposed responsibilities of the Board and Senior Management in the Guide.

		However, we seek further guidance on the specific requirements by the Board as part of addressing AI risk within the FI's risk appetite framework under para 2.5b of the proposed AIRG to include the setting of appropriate qualitative statements and quantitative measures or limits. This technical process is usually handled at the functional risk management forums in the FI. **Question 4:** We agree with 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. **Question 5:** We agree with the proposed approach for the identification of AI. However, we seek further clarity on the specific criteria and guidelines for identifying AI in the context of third-party service providers to ensure consistent and effective identification of AI across all relevant business and functional areas, including those involving external partners. We encourage MAS to provide additional guidance on how FIs should apply these criteria when engaging third-party service providers to maintain alignment with the Guidelines **Question 6:** We agree with the proposal for FIs to establish and maintain an accurate and up-to-date inventory of all AI usage. However, we recognize the challenges involved in obtaining timely and comprehensive information related to third-party AI models, particularly when changes are made to existing models or services. Additionally, we seek MAS's clarity over "up-to-date inventory of all AI usage", if MAS' intention is to cover for all, including pilots and Proof-of-Concept's (POCs). If so, it should be clear and align across the industry where clarity is provided on the expectations for the need to balance innovation and risk controls. **Question 7:** We agreed with the proposed risk dimensions of Impact, Complexity and Reliance that covers the AI risk materiality. **Question 8:** We are cognizant of the challenges relating to the proposed standards, processes and controls applied for Gen AI use cases. 1. "Black-Box" (proprietary information) of open sourced LLMs making the evaluation, re-training, as well as the conduct of Fairness and Transparency assessment on Gen AI use cases difficult. 2. There were concerns revolving around the reproducibility and traceability of Gen AI use case. What would be the expectation from MAS on this?

		Our suggestion is for MAS to consider appropriate industry accreditation of third-party providers which would aid FIs to assess and onboard the service provider. Referring MAS Outsourcing Guide, January 2025, Annex 2, which provides examples of outsourcing agreements that would not be deemed material outsourcing for FIs. How would this be read in conjunction with the MAS AI Risk Management Guide under the Third Party AI Life cycle control? **Question 9:** Given that the definition and recommended controls in the MAS AI Risk Management Guidelines are based on the current state of technology, how does MAS envision maintaining the relevance and effectiveness of these Guidelines as AI technology evolves and advances? What are MAS's expectations regarding balancing robust risk management controls with fostering innovation in the context of future AI developments? Additionally, are there any aspects of AI risk management related to these advanced AI techniques that have not been sufficiently addressed in the proposed Guidelines? We seek MAS guidance under what circumstances testing at the model level alone is sufficient, and whether use cases that share the same model but differ only by prompts, or applications that allow users to input dynamic prompts, require separate use case-level testing. While autonomy is considered, the degree of autonomy is not defined or classified clearly, i.e., what is considered autonomous, particularly in scenarios like multi-agent frameworks where both autonomous processes and human-in-the-loop mechanisms coexist. **Question 10:** A phase approach could be considered to establish and operationalise the prescribed recommendations from the Guide, (e.g., transition over Phases). Phase 1 – We opined that 12 months for the transition to enhance the respective risk framework and policies. Phase 2 – Additional time (approx. 12 months) should be considered for the operationalization of the control standards including processes, which would depend on the magnitude/scale of the organisation, especially when rolling out an enterprise-wide governance across subsidiaries, overseas branches and third-party service providers.
56	Palo Alto Networks	**Question 1:** Palo Alto Networks thanks the Monetary Authority of Singapore for the opportunity to comment on the Proposed Guidelines on Artificial Intelligence Risk Management for Financial Institutions. With financial institutions sitting at the center of the digital economy, the imperative for the sector should be clear: simultaneously embrace AI for cybersecurity, and cybersecurity for AI.

		We offer the following key recommendations to ensure that security and innovation are foundational components of Singapore's AI Risk Management Guidelines for financial services, and as outlined in greater detail below: PAN Recommendation: Support the risk-based and proportionate application . Recommend that minimum security mandates apply to all FIs using AI as an integrated part of their business processes, regardless of size. Mandate Foundational Policies: all FIs should institute basic policies for AI adoption built on Secure AI by Design imperatives, which include securing the use of external AI tools, securing underlying AI infrastructure and data, safely building and deploying AI applications, and monitoring and controlling AI agents. These guidelines should empower organizations to combat "Shadow AI" and disallow unsanctioned uses or applications of AI; prohibit inputting PII data into public AI tools, monitor for emerging attack vectors such as prompt injection, model extraction, and data exfiltration; implement mechanisms to block unsafe or risky models and patterns in production; and continuously assess security, safety, and compliance risks across infrastructure, applications, models, and datasets - both in the supply chain and in runtime - through automated posture management, model scanning, and continuous adversarial testing. Proportionality in Technical Controls: The most stringent technical controls, such as AI Red Teaming, should apply to clearly defined high-risk use cases, like those AI deployments where the system would be making material decisions that impact individuals' access to critical services.. Platform Approach: Proportionality should be achieved by selecting scalable security platforms, which allow FIs to centrally manage and calibrate granular controls efficiently, rather than relying on fragmented point solutions as well as have continuous monitoring and automated remediation. **Question 2:** PAN Recommendation: Endorse the comprehensive scope that includes Generative AI and AI agents. Generative AI Risks: The scope should capture GenAI risks, including prompt injection, data leakage, and malicious content from responses. AI Access Security and AI Runtime Security are necessary for mitigating these newly amplified attack vectors. AI Agent Risks: The inclusion of AI agents is critical, given they introduce more significant risks due to their greater autonomy and ability to access tools and execute commands across other parts of a network. AI Agent Security directly addresses these risks, such as memory manipulation and tool misuse.

		Ecosystem Components: The scope should explicitly cover all components of the AI ecosystem, including apps, plugins, agents, models, datasets, data pipelines, and infrastructure. **Question 3:** PAN Recommendation: Emphasize that effective Board and senior management oversight depends on access to continuous, automated technical risk intelligence, and is a critical aspect of an enterprises' overall risk management processes. Promote Zero Trust: The Board should mandate Zero Trust Architecture (ZTA) as the foundational security system to secure distributed AI data pipelines and infrastructure. This ensures that FIs can meet their responsibility to ensure explicit and continuous verification of identity and permissions for access to sensitive AI components. Continuous Posture Visibility: Senior management should require continuous monitoring and remediation of the FI's AI security posture, achievable through AI focussed Security Posture Management. This provides the necessary technical insights for updating existing risk management frameworks (e.g., technology and cybersecurity) to address AI risks. **Question 4:** PAN Recommendation: Support the establishment of a dedicated committee for material risk exposure. Recommend that overall AI risk exposure assessment incorporate quantifiable, technical risk dimensions derived from continuous monitoring. Input for Assessment: The assessment of overall material risk should aggregate and correlate findings from: AI Misconfiguration Risk: Data revealing widespread misconfigurations (e.g., excessive permissions) and cloud risks across AI applications, data, and pipelines. GenAI Usage Risk: Metrics from AI inventories and technical capabilities, like AI Access Security tools, to map usage that can detail the percentage of GenAI applications that are unsanctioned or where employees have attempted sensitive data leakage. Third-Party Dependence: Quantification of concentration risks from over-reliance on a few dominant GenAI providers. **Question 5:** PAN Recommendation: Recommend that "robust systems" should include AI Access Security and AI security posture management (AI-SPM) to automate the discovery process and prevent shadow AI. Automated Discovery: These systems provide comprehensive visibility into shadow AI and the broader AI ecosystem.

		Extensive Cataloging: The technology should leverage dictionaries, offering a catalog of GenAI applications, alongside deep visibility into GenAI-specific attributes by risk. Preventing Shadow AI: Automated discovery is essential for mitigating model sprawl and shadow AI models, which lack governance and increase the attack surface. **Question 6:** PAN Recommendation: Recommend AI-SPM capabilities for automated maintenance of the AI inventory to ensure it is continuous and technically accurate. Continuous Inventory: An inventory should be continuously updated to keep pace with rapid DevOps deployments and cloud-native application development, countering model sprawl. AI-SPM provides a live model inventory. Mandatory Technical Attributes: The inventory should capture technical attributes essential for security and governance, including: AI Lineage: Identifying and tracing the lineage of AI components and data sources used in applications. Model Access Governance: A visual mapping of who has access to deployed AI models and associated resources. Training Data Classification: Identifying whether models are trained or fine-tuned on sensitive data. **Question 7:** PAN Recommendation: Endorse the proposed dimensions. Recommend the inclusion of a quantifiable, security-specific dimension. Additional Dimension: Security Posture Risk / Attack Surface Exposure: This dimension measures the technical integrity and configuration security of the AI environment, which directly contributes to the inherent risk materiality. Metrics: It should track risks identified by AI-SPM, such as excessive permissions, sensitive data exposure, and platform/access misconfigurations. Rationale: Misconfigurations in data pipelines and deployment infrastructure can introduce significant security and compliance risks, allowing attackers to exploit vulnerabilities. Assessing this technical posture provides proactive metrics to evaluate the true potential impact of a failure. Assessing and Monitoring Residual Risk: The impact, complexity, reliance are the inherent risk factors. Security Posture is effectively measuring control (configuration) effectiveness; i.e. a residual risk view.

		Impact (Materiality of Impact): Evaluates the potential damage or consequence to the organization, customers, or financial system if the technology fails, acts unexpectedly, or produces biased outputs. This includes financial loss, legal or regulatory breaches, and harm to customers. Complexity: Measures the intricacy of the technology, such as the sophistication of AI models (e.g., Generative AI), the opacity of "black box" algorithms, or the level of autonomy in AI agents. Reliance (Dependence): Measures the extent to which the institution relies on the technology for its core functions. High reliance increases vulnerability to operational disruption or "over-reliance," where human decision-making is compromised by, or delegated to, faulty technology. **Question 8:** PAN Recommendation: Recommend specific technology platforms for enforcing standards in high-risk areas like third-party risk management, evaluation, and technology security. Key Area : Third-Party AI & GenAI Risk PAN Recommendations : Recommend SASE/AI Access Security Rationale : Necessary for inline enforcement to prevent the input of confidential data into public AI tools. Provides real-time, context-aware LLM-powered data classification. Also blocks malicious URLs and malware in GenAI responses. Require AI Model Security (model scanning) to check third-party models for vulnerabilities before deployment. Key Area : Evaluation and Testing PAN Recommendations : Recommend AI Red Teaming Rationale : FIs should implement adversarial testing. AI Red Teaming provides continuous assessment using a dynamic AI agent to adapt and test against attack scenarios across various techniques. Key Area : Technology & Cybersecurity PAN Recommendations : Recommend AI Runtime Security & ZTA Rationale : Require controls against threats to deployed models. AI Runtime Security protects against prompt injections and implements custom topic guardrails. Require network segmentation and micro-segmentation for deployment environments. Key Area : AI Agents PAN Recommendations : Recommend AI Agent Security Rationale : Must secure the deployment and actions of autonomous agents by blocking threats like identity impersonation, memory manipulation, and tool misuse. Key Area : Reproducibility & Auditability PAN Recommendations : Recommend Model Security & Logging Rationale : Require using AI models that are duly scanned for 3rd party risks and tracked using stringent inventory controls. Ensure audit trails are maintained for all agent

		interactions and runtime security events. Logging and monitoring of events and incidents in a central location continuous monitoring. **Question 9:** PAN Recommendation: Promote Zero Trust Architecture (ZTA): ZTA should be promoted as a best practice and a foundational requirement for securing distributed AI data and infrastructure. The focus should shift from managing a network perimeter to continuous verification and least privilege access across all AI components (data lakes, models, compute). Redefining Capacity via Autonomous Security: The Guidelines should recognize that human capacity is inadequate against the speed of AI-accelerated threats. Capacity and capability should be measured by the FI's level of automation. FIs should prioritize investment in Autonomous Security Platforms enabling AI enabled detection and automated remediation. This ensures the FI's security operations can perform real-time detection, prevention, and resolution of incidents, which is the only way to meet high-velocity threats in the financial sector. Continuous Evolution: The AI threat landscape is evolving rapidly, establishing practices and solutions to continuously monitor evolving threats and update relevant processes and controls. Actionable Intelligence: Leverage platforms and solutions that make use of up-to-date, actionable intelligence related to new threats and attack techniques in monitoring AI security posture, AI testing and evaluation. **Question 10:** PAN Recommendation: Agree that 12 months is an appropriate transition period, provided FIs focus on implementing platform-based solutions that expedite compliance for the most material risks. Focus on Integrated Platforms: FIs should prioritize deploying integrated platforms like SASE and AI-SPM immediately, as these solutions offer unified security, simplifying management and accelerating the adoption of necessary controls (DLP, access control, posture management). Time-to-Compliance Acceleration: Integrated solutions are essential because they reduce the time needed to gain comprehensive visibility and implement technical guardrails across the organization, making the one-year timeframe achievable, even for complex ZTA implementation.
57	Ping Identity	**Question 1:** We endorse the principle of proportionality but advocate that the threshold for regulatory scrutiny must pivot from the institution's size to the agent's autonomy and actionability. In the emerging "Agentic Economy," a small financial institution deploying fully autonomous "Digital Workers" capable of executing transactions faces a risk profile

		comparable to a global bank. We recommend a governance framework that categorizes AI not merely by complexity, but by Identity Type—distinguishing between "Personal Agents" (unmanaged, low-trust) and "Digital Workers" (managed, high-trust). Furthermore, proportionality must account for the interface: Computer-Using Agents (CUAs) that manipulate GUIs present inherently higher audit and security risks than API-driven agents and warrant rigorous oversight regardless of the deploying entity's scale. **Question 2:** To future-proof the regulatory framework, the scope must expand beyond "models" to encompass the integration protocols that enable autonomy, specifically the Model Context Protocol (MCP) and Agent-to-Agent (A2A) interfaces. The definition of an AI system should explicitly include the Non-Human Identity (NHI) assigned to the agent. An AI model is merely a calculation engine; an AI Agent is an active digital identity capable of tool execution and data retrieval. To prevent "Shadow AI," the scope must include agents that dynamically generate sub-agents to perform tasks, ensuring that these transient identities and their recursive delegation chains remain within the regulatory perimeter. **Question 3:** Governance at the Board level must shift from model approval to Identity Sovereignty. We propose that Senior Management be held accountable for enforcing a "Delegation, Not Impersonation" mandate. Policies must strictly prohibit AI agents from sharing or "borrowing" human credentials, which obscures audit trails and liability. Instead, leadership must institute architectures where agents operate via Authenticated Delegation—possessing distinct, verifiable identities that act on behalf of a human but remain clearly distinguishable in forensic logs. Additionally, Management must oversee the implementation of Human-in-the-Loop (HITL) verification for high-impact financial decisions to ensure non-repudiation. **Question 4:** Given that Agentic AI operates across traditional silos—accessing APIs, databases, and external SaaS tools simultaneously—this committee must elevate Identity and Access Management (IAM) leadership to a strategic risk function. Traditional security models focus on perimeter defense; however, autonomous agents act as insiders within the trust boundary. This committee must oversee the implementation of a Unified Control Plane that governs agent entitlements, ensuring that an agent's ability to "act" (e.g., read PII, execute trades) is strictly governed by Zero Trust principles of least privilege, identical to the rigor applied to human employees. **Question 5:** Static registries are insufficient for high-velocity agentic environments. We recommend mandating Dynamic Client Registration (DCR) standards for AI agents to ensure that every agent is assigned a unique, cryptographically verifiable identity upon instantiation. Identification systems should utilize metadata standards, such as "Agent Cards," which declare an agent's capabilities, authentication requirements, and authorized data modalities. This ensures FIs can instantly identify not just which model is active, but what tools it is technically capable of accessing, thereby closing the visibility gap on "rogue" or "shadow" agents.

		Question 6: The AI inventory must evolve into a dynamic Identity Registry that tracks the chain of delegation. It is critical to record not only the agent's existence and provenance but the human sponsor who authorized it and the specific scopes granted (e.g., read-only vs. execute transaction). Given the ephemeral nature of some agents, the inventory system must capture transient lifecycles to prevent "ghost agents"—identities that persist and retain high-privilege access after their utility has expired. This supports a robust "Joiner-Mover-Leaver" lifecycle management process for non-human identities. **Question 7:** We propose adding "Autonomy Level" and "Actionability" as critical risk dimensions. The risk profile shifts fundamentally when an AI transitions from providing advice (read-only) to executing workflows (read-write). Furthermore, risk assessments must evaluate the interaction interface. Agents operating via Computer-Using interfaces (CUAs) that simulate human inputs on a screen pose significantly higher security risks than API-based agents because they bypass standard API security controls and are brittle to audit. Such agents should automatically trigger a higher risk materiality classification. **Question 8:** Lifecycle controls must prioritize Identity Lifecycle Management (ILM) for agents with the same rigor used for human employees. Authentication & Oversight: Controls should consider standards such as Client-Initiated Backchannel Authentication (CIBA) for high-sensitivity actions. This ensures that when an agent acts on behalf of a user, the user receives an out-of-band notification to explicitly approve the transaction, effectively digitizing the "Human-in-the-Loop" requirement. Authorization: FIs must enforce Just-in-Time (JIT) access, where agents are granted ephemeral credentials only for the duration of a specific task, minimizing the attack surface if an agent is compromised. **Question 9:** The Guidelines should explicitly address Interoperability and Agent-to-Agent (A2A) Trust. As FIs deploy multi-agent systems where agents collaborate (e.g., an internal "Digital Worker" delegating a task to an external "Payment Agent"), regulations must specify how trust and context are securely preserved across these boundaries. We also recommend guidance on "Secretless" Identity architectures. To mitigate credential theft, guidelines should encourage systems where agents do not store static secrets (like long-lived API keys) but rely on dynamic tokens brokered by a secure gateway. **Question 10:** Given that a majority of organizations are already piloting autonomous agents, a 12-month transition may leave FIs vulnerable to unregulated "Shadow AI" risks. We suggest a phased approach where Identity Governance—specifically the discovery and registration of all active agents—is prioritized immediately. Establishing visibility into what non-human identities are currently operating on the network is a prerequisite for managing the more complex risks of bias, explainability, and adversarial defense.
58	PricewaterhouseCoopers	**Question 1:** No further input.

	Risk Services Pte Ltd	**Question 2:** We are broadly supportive of the proposed scope of AI. Addressing Agentic AI Systems We recommend MAS to future-proof the scope of AI by explicitly including an Agentic AI System as a distinct category capable of autonomous planning, tool use, long-horizon task execution, and self-improvement. Currently, as defined in Section 3.3 (b), an AI system can comprise one or more models and other machine-based component. It is not explicitly clear if this comprises a system of Agents or an Agentic system. The current framework, while relevant and useful for traditional and Generative AI, requires explicit extension for this emerging class. **Question 3:** We broadly concur with the roles and responsibilities of both the Board and Senior Management in AI risk management oversight. Roles and Responsibilities in the Event of a Public Crisis We kindly request MAS to consider detailing the Board's and SM's role in the event of a public crisis. AI-related incidents are inevitable when usage becomes prevalent and if it does, trust in the financial services ecosystem will be impacted. Trust is a cornerstone of the financial services system and hence we recommend for crisis management responsibility to be spelt out more explicitly under the proposed responsibilities of Board and Senior Management. Communications with main stakeholders, including investors, regulators and the public is key to maintaining trust and managing reputational risk. Boards should lead a thorough review of the event and the effectiveness of the response after managing the crisis. They should then advocate for necessary changes to policies and strategies to better prepare for future crises. **Question 4:** A dedicated cross-functional committee to oversee AI risk is a sound approach to ensuring enterprise-wide AI risk oversight. To operationalise this, we kindly recommend several considerations to be taken into account to provide more latitude to FIs to manage material AI risk. * Global/regional reliance: Clarify the extent to which FIs (especially branches or foreign subsidiaries of foreign entities) may rely on global or regional AI committees, especially where such forums sit outside Singapore. Allowing local management to leverage these structures while retaining local accountability would support practical implementation. * Flexibility in committee structure: FIs should have the option to either establish a dedicated AI risk committee or adapt existing committees (e.g., Technology Risk Management & Cybersecurity, Model Risk, Risk Management) by expanding their terms of reference to cover material AI risk, achieving the intent of this proposal. * Definition of “cross functional”: While FIs should retain flexibility in determining committee composition, it is recommended for business representation to be more

		explicitly indicated or encouraged. Business teams are closest to customer facing AI use, which carries inherently higher risk and is often underrepresented compared to technology, data, model, and risk functions in such forums. Addressing Agentic AI Systems and Use Cases We kindly recommend the charter mandate at least one member with certified expertise in Agentic AI safety. **Question 5:** We kindly recommend MAS to adopt a short, industry-consistent "AI Identification Decision Tree" that every FI must run whenever a new system is built or procured. We kindly recommend MAS be more explicit in mentioning all AI usage/exposures across all relevant business and functional areas. **Question 6:** Calibrated, risk-based inventory structure We kindly recommend a calibrated, risk-based inventory structure. Only material or high-risk AI use cases should require the full set of attributes, while low-risk cases should complete only a minimal subset to maintain operational feasibility and reduce administrative burden, aligned with MAS's intent of a proportionate, risk-based approach. Low-risk cases should still remain in the inventory to support periodic risk reviews and ongoing monitoring, ensuring they are reassessed if their risk level changes from a low-risk to a high-risk one. Alignment with AI Verify We kindly recommend MAS to consider the AI Verify Inventory Schema as the industry reference, including mandatory attributes for Agentic capability, autonomy level, tool access and safety guardrails. **Question 7:** Risk dimension to focus on business use of AI and linkage to critical business services We kindly recommend adding business use of AI, specifically its linkage to critical business services and functions (aligned with MAS Business Continuity Management Guidelines), as a separate risk dimension. Focusing too heavily on AI complexity may obscure what truly matters: a complex AI used in non-critical services should not automatically be deemed material, while AI embedded in critical services warrants closer attention. Addressing Frontier Models, Systems and Use Cases The current triad (impact, complexity, reliance) is sound but may not be adequate for Frontier models. We kindly recommend adding the following dimensions to the materiality assessment framework: * Autonomy level * Reversibility of decisions/outcomes * Data sensitivity & personal data volume * Presence of feedback loops / self-reinforcement risk * Cross-border data or model flows * Explainability / interpretability score

		* Degree of Third-Party / Generative AI dependency These additional dimensions are particularly critical for Generative AI and Agentic systems. **Question 8:** Third-Party AI Management We acknowledge that FIs should remain fully accountable for Third-Party AI and customer outcomes. While MAS expects appropriate reviews of Third-Party and open-source AI models, including supply chain risk assessments and validation of model provenance, data integrity, and known vulnerabilities (para 4.11, 4.11c), this is challenging in practice, especially for Generative AI because: * Large technology providers may not provide sufficient transparency. * Technology vendors have limited commercial incentive to share detailed information unless demanded collectively. * Expertise required to assess Third Parties may be scarce particularly for smaller FIs. To address this: Would MAS consider explicit guidance on minimum standards that third-party providers should be assessed on AI. Cyber testing (Paras 4.14b and 4.22b) As per good cyber security testing practices, we believe threat modelling would be good practice to enable threat events identification, consistent risk quantification and control design / implementation. Change Management (Para 4.25) We seek clarification from MAS as to whether risk and impact assessments should be performed as part of the change process, including identification of applicable dependencies as detailed in the MAS Change Advisory. Addressing Agentic AI Systems and Use Cases Given the nature of Agentic AI, certain more sophisticated security and controls may need to be considered (all field-tested in the AI Verify Agentic AI Testing Framework): * Pre-deployment Agent Red-Teaming & Scaffold Safety Assessment * Run-time Goal & Instruction Hardening - Mandatory circuit-breakers and stop-button efficacy testing - Continuous Autonomy Monitoring with tiered escalation - Tool-access whitelisting and privilege separation - Memory & state sanitisation controls. **Question 9:** Roles of 3rd and 4th Lines of Defence We believe the assurance functions are particularly relevant as part of the AI risk framework as: * AI risks are new, and 2nd line risk safeguards will be maturing; * 'End-to-end' holistic lifecycle controls span multiple departments in FIs, where hand-off risks are higher and assurance functions have been strong in calling these out.

		To this end, the guidelines would benefit from a build out of the 3rd line (Internal Audit) and 4th Line (External Assurance) roles to ensure these safeguards keep up with developments in Line 1 and Line 2. Singapore has had significant success on AML/CDD with the 'AAPG' group collaboration across Internal and External Audit to lift standards and transparency on the work done by 3rd/4th line. We believe it is beneficial for the AI guidelines to consider a similar working group now (when AI risk frameworks are at the early stages of being developed), given the integral roles 3rd and 4th lines play, rather than these coming in much later. The consultation paper discusses governance roles. One learning from the AAPG is that cohesion between 'checkers' is critical to prevent fatigue. As such, we believe the governance layer (e.g. Audit Committee) should have line of sight of the scope of checks over AI planned by the different Lines (Line 1.5, 2, 3, External Audit). In the same vein, the governance layer will also need to have a line of sight over the quality of internal and external audit tools (including AI) used in providing assurance for the Board. We kindly request MAS to take this into consideration. The larger audit firms also have global reach/connectivity. AI Risks and safeguards are evolving rapidly, and larger audit firms can bring to the working group up to date practices on global trends and best practices in AI safeguard frameworks. Addressing Critical Gaps in Agentic AI Systems and Use Cases We kindly highlight four priority gaps that warrant specific attention in MAS guidance: * Governance for Agentic and multi-agent systems – Frameworks to manage autonomy, decision rights and accountability across Interacting Agents. * Human-agent interaction standards – Clear expectations for human competency, acceptable latency, escalation protocols and human oversight in mixed human – AI operations. * Concentration and systemic risk – Assessment and mitigation of risks arising from widespread dependence on common Agentic platforms or shared models. * Crisis management risk – AI-related incidents especially those caused by AI Agents are not known to us and these incidents may create a larger fall-out than expected. In addition to have kill-switches, we recommend that the risk of AI agents going rogue be explicitly considered. Addressing Frontier Models, Systems and Use Cases We kindly recommend the Board receive an annual briefing on frontier risks, specifically Agentic and Autonomous Decision Systems. Consideration of external risks/threats The Paper's framework is comprehensive, however it focuses on risks from a FI's internal use of AI. This could be interpreted as 'not applicable' by FIs that do not use AI/Generative AI/Agentic AI models. However, there are also increasing AI risks from external 'threat actors'. For example, on cybersecurity and scams, the risk landscape is rapidly changing due to threat actors using

		these technologies (MAS/TCRS/2024/05: Cyber Risks Associated with Generative Artificial Intelligence). We feel the Paper would benefit by explicitly encompassing the external changing risks/threats, so that FIs keep abreast of these too as well as internal AI risks, including if the entities do not use AI/Generative AI/Agentic AI models themselves. **Question 10:** Staggered transition We agree with the minimum 12-month transition period but recommend MAS to consider a staggered and extended approach for some aspects of the Guidelines (for example, evaluation, testing and explainability of Agentic systems) to allow for further industry readiness. A proposed progressive roadmap for consideration is as follows: * Months 0–12: Full compliance for existing traditional & Generative AI * Months 12–18: Inventory & materiality for all Agentic pilots * Months 18–24: Full Agentic lifecycle controls for production systems
59	QCP Trading Pte Ltd	**Question 1:** QCP supports the proportionate application of the guidelines, taking into consideration the risk and scale of the organization deploying AI. However, QCP seeks further clarification on the definition of proportionate application and how FIs may determine an appropriate proportionate application of the Guidelines. We also seek clarification on the requirements for regular compliance checks against the Guidelines (as specified in item “d”, Paragraph 5 of Annex) as well as the geographic scope of the guidelines if they only apply to AI use cases supporting processes for clients in Singapore or if it includes overseas clients but managed/advised from Singapore. To further support implementation of AI across the industry, MAS may wish to consider providing illustrative supervisory expectations for the different FI size and operational complexity (smaller setups vs complex regional groups) and clarify the evolvement of AI usage as it matures/scales over time in terms of proportionality. **Question 2:** QCP suggests that the scope provide more clarity on whether it extends to the increasing trend of vendors (e.g. Salesforce) implementing AI into their systems and whether such cases should be regulated under this guideline or the MAS Outsourcing Guidelines. There could be cases where the vendor may also implement AI without informing the FIs and cases where it is unfeasible for the FI to have complete oversight over every 3rd party software. A risk-based approach to focus on material software vendors identified already via the MAS Outsourcing Guidelines could be considered to provide more clarity in this regard. Also, where AI embedded within third-party vendor platforms is considered in-scope, we seek clarification on whether such AI is expected to be included in the AI inventory as set

		out in section 3.4 of the proposed Guidelines. In such cases, we would also like to clarify whether it would be reasonable to record such AI in the inventory at a higher level (e.g. at an awareness or use-case level) as detailed model-level information may not be available or accessible to us. QCP also suggests leveraging existing international resources (e.g. OECD AI) to avoid fragmentation across the industry. QCP seeks clarification on what is considered "materially dependent" AI and how the threshold should be defined. The material impact assessment should be based on measurable impact on clients and services instead of technology type, AI or otherwise. In the same vein, does MAS expect the whole application or platform to fall within scope, or only the AI-enabled parts and use cases? Should there be retrospective identification and reassessment when there are such changes? **Question 3:** As AI technologies become increasingly integrated into financial services, it is essential that the highest governing body ensure AI adoption is aligned with the institution's strategic objectives, risk appetite, and regulatory obligations. We support MAS' proposal for the Board to oversee the overarching governance framework and endorse key AI-related policies. **Question 4:** We support the risk-based approach for the establishment of a dedicated AI committee. However, we would also propose an allowance for smaller FIs to subsume said committee into the larger/enterprise risk committee instead. With that said, does MAS expect the assessment to be qualitative or quantitative, primarily? Or a combination of both? How should interdependencies between AI use cases be factored into the assessment? Are there any common supervisory indicators that may suggest AI risk exposure is material (e.g., use of AI in customer-facing decisions or regulatory processes?) **Question 5:** In line with the risk-based approach by the MAS, we support the proposal for FIs to have the flexibility to establish the definitions and criteria to identify AI across the business. Additionally, does MAS have any expectations on FIs to manage the risks of AI usage that arise from decentralized experimentations or third-party tools where usage evolves quickly? **Question 6:** We seek clarifications on whether the MAS will prescribe a review frequency and the expectations on maintaining the inventory of AI usage. Drawing parallels with the TRMG and Outsourcing Guidelines, QCP suggests that the MAS adopt a similar approach in inventory where the focus is on material AI usage. As more

		software products/services embedded AI within them, it creates an avenue of concern on how to identify where AI is integrated within the product. Instead of tracking all instances of AI usage, the focus should be on creating an inventory of AI usage with material impact. Further, QCP seeks clarification on whether MAS intends to prescribe a standard template for the AI inventory or register, like the approach taken for the Outsourcing Register, and whether such an AI inventory or Register would be subject to audit. **Question 7:** QCP suggests that complexity may be redundant as it will already be covered under the impact assessment. It is noted that many global frameworks (ISO 23894:2023 – severity, probability , EU AI Act – use case, impact)) have moved away from using complexity as a factor in assessing materiality as it does not reliably reflect actual risk, QCP suggests that the MAS consider focusing on the use case and impact instead. To further clarify, if the business use case remains unchanged, does the MAS expect risk assessments to be updated when there are material changes to data pipelines, model behaviors, or operational environments? **Question 8:** QCP seeks clarity on the transparency and fairness requirements under Paragraph 4.9. We seek more concrete direction on what fairness testing should encompass and how FIs can identify and monitor protected attributes. The shared responsibility model mentioned in Paragraph 4.11 implies that FI bear end-to-end responsibility even when relying on external providers, QCP suggest that MAS consider accepting assurance reports from providers, like the approach with cloud services where third party suppliers should be accountable for aspects under their control. We also suggest that the MAS specify the expectations for suppliers such as types of attestations or documentations that can be accepted across the industry. QCP suggests adopting a risk-materiality test similar to the EU AI Act to ensure enhanced oversight for material AI components that may affect FI decision making or customer outcomes. We agree with the MAS' focus on reproducibility and monitoring but suggest shifting the focus from exact result replication to look at auditability of methodologies (processes, templates, etc.) to accommodate behavioral drift in generative models. In addition to the above observations and suggestions, QCP seeks clarification on the following aspects of the proposed standards in support of consistent and proportionate implementation. 1. Are specific life cycle controls mandatory or optional for low-risk AI tools where it acts as an assistive tool or productivity enhancer? 2. Would there be illustrative examples MAS would be providing for mapping of risks tiers for expected life cycle controls for Gen AI? 3. In scenarios where there is limited transparency into training data, model architecture, how do we demonstrate compliance with requirements? In such cases where full life cycle

		transparency is not practicable, what kind of compensating controls can FIs implement risk assessment to be deemed sufficient? 4. AI models could be constantly updated. How should FIs balance the depth and independence of review against the speed of development? 5. For AI models/agents with deemed higher autonomy, are there any expectations around governance and testing of a kill-switch mechanism? **Question 9:** While the Guidelines are comprehensive, QCP seeks clarifications on the following: 1. Where FIs rely on third-party AI solutions and provide client data as inputs, we would like to seek clarification on the regulatory implications for the FI itself in the event of a data breach at the third-party. Would this be a reportable incident to PDPC and MAS at the FI level? 2. Additionally, we would like to seek clarity on the incident management and reporting format for AI-specific incidents. Would the same FI incident reporting process apply here? 3. For FIs operating across multiple jurisdictions, would MAS expect SG regulated entities to adopt stricter standards if group level frameworks are less developed? Or in equivalence would be acceptable as well. 4. In scenarios where staff independently adopt AI tools, would MAS expect FIs to actively detect such usage? Or would education be sufficient? **Question 10:** We propose a longer transition period of 18 months at least, in view of the time required to build the necessary capabilities, address gaps in expertise, and allocate sufficient resources to implement the proposed guidelines effectively.
60	Rapyd Holdings Pte Ltd	**Question 1:** We agree that AI governance should be based on risk, lifecycles and be proportionate. **Question 2:** We support the proposed scope generally, though we recommend greater clarity on what constitutes “AI use cases, systems and models” for the application of the Guidelines. In particular, deterministic, rule-based systems or low-risk automations (such as scripted chatbots, simple workflow tools or automations) should be clearly excluded or treated as out-of-scope. Clear distinctions would help FIs focus governance efforts on higher-risk AI use cases rather than benign automation. **Question 3:** We generally agree with the proposed allocation of responsibilities between the Board and senior management, for the Board to set overall risk appetite, governance expectations, and oversight, whilst senior management is responsible for operationalising AI risk management frameworks. However, we encourage MAS to continue emphasising proportionality, particularly for boards of fintechs where AI may support, rather than drive, core regulated activities. **Question 4:** We support the proposal generally, provided it is triggered only where the overall AI risk exposure of an FI is deemed material. The key issue is how “materiality” is assessed. If FIs

		are allowed to self-assess materiality using clear, risk-based criteria, this requirement would not be unduly onerous. **Question 5:** We agree with the need for clear definitions, criteria, and processes to ensure consistent identification of AI usage across the organisation, particularly if implemented in a pragmatic manner. Overly broad definitions may unintentionally capture low-risk tools and create unnecessary compliance costs. We would appreciate if further guidance or examples illustrating borderline cases could be provided, to assist in achieving consistent and efficient identification. **Question 6:** As AI usage governance should be based on risk and proportionality, FIs should be permitted to integrate AI inventories into existing system or model inventories, rather than creating standalone repositories. **Question 7:** We support the proposed criteria for risk dimensions; we would encourage greater clarity on the treatment of inhouse Small Language Models ("SLMs") vs third party large language models (LLMs), as SLMs are less complex and lower risk. **Question 8:** We support the proposed lifecycle standards and controls and agree that they should be applied in a proportionate, risk-based manner. We note in particular the importance of data management and privacy safeguards. However, we recommend that the Guidelines avoid introducing restrictions on data usage or data uploads to AI systems beyond what is already required under existing data protection and confidentiality, and security laws. As this is an evolving area, maintaining flexibility in how FIs use AI is important to ensure continued innovation and responsible adoption of AI. **Question 9:** We hope that MAS may consider providing additional guidance on cross-border AI usage, particularly where global AI models or platforms are used across multiple jurisdictions. Alignment with international AI governance frameworks would help multinational FIs manage compliance consistently. We also encourage MAS to continue engaging industry participants as AI technologies evolve, to ensure the Guidelines remain practical and innovation-enabling. **Question 10:** We support a transition period of at least 12 months. This provides a reasonable timeframe for FIs to conduct gap assessments, update policies, and implement governance structures. For FIs with complex global operations or reliance on third-party AI vendors, a phased implementation approach focusing first on higher-risk AI use cases would be particularly helpful.
61	RHB Bank Berhad	**Question 1:** To support consistent supervisory outcomes, could MAS be prescriptive on the quantitative and qualitative indicators for "materially dependent" (e.g. direct contribution

	Singapore Branch	to core business processes, number or % of customers affected), which are key concepts used to determine if AI is an integrated part of our business processes and therefore, is subject to more extensive requirements.] **Question 2:** For practical implementation, we suggest MAS to include in the final Guidelines or accompanying FAQ, a short taxonomy of illustrative use cases, distinguishing (a) AI systems that would generally be treated as material (e.g. customer-facing decisioning, financial crime, risk measurement) from (b) low-risk productivity tools (e.g. office copilots used solely for drafting internal documents) where only basic safeguards would be expected **Question 3:** We are of the view that Senior management's responsibility as drafted in the Guidelines is adequate, and there will not be a need for Board (or delegated committee) to have oversight at the operational level, which is Senior management's responsibility. The Board, or a committee delegated by it, is responsible for: a. approving the overall governance approach for AI risk management, including key frameworks, structures, policies and procedures designed to assess and manage the FI's AI risk on an ongoing basis; b. ensuring that AI risks, where material, are explicitly addressed within the FI's risk appetite framework, including the setting of appropriate qualitative statements and quantitative measures or limits; 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. Senior management is responsible for: a. ensuring the effective implementation of AI-related risk management policies and procedures across the FI, consistent with the FI's risk appetite; b. regularly reviewing the effectiveness of the AI-related risk management 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. ensuring robust mechanisms for coordination and accountability for AI-related risk management are established and maintained across the FI; d. establishing an internal escalation process for managing material AI risks and exceptions, such as incidents or breaches of risk thresholds, and ensuring appropriate and timely actions are taken; e. updating the Board on material AI risk issues in a timely manner; and f. 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.

		Question 4: To avoid inconsistency in the establishment and scope of such committees, we suggest that MAS: - Clarify, in guidance or examples, the types of AI usage where MAS would normally expect a cross-functional committee (e.g. multi-business line use, systemic importance, significant reliance on external AI providers), in line with the expectation of the guidance on risk materiality assessment set out in paragraphs 3.8 to 3.11 of the Guidelines. **Question 5:** Nil **Question 6:** Nil **Question 7:** Nil **Question 8:** Nil **Question 9:** Request MAS to clarify expectations where AI systems are developed, hosted or operated by group entities or Head office, but used by a Singapore-regulated FI. **Question 10:** Implementation timelines will vary depending on the scale and complexity of AI usage and third-party dependencies. We suggest: - FIs should prioritise applying the detailed expectations to AI systems assessed as high-risk within the 12-month period, while permitting a documented, risk-based phasing for lower-risk use cases. - Indicate that MAS will take into account an FI's documented implementation plan, governance arrangements and good-faith progress when assessing adherence during and shortly after the transition period, particularly where external dependencies (e.g. vendors) affect timelines.
62	Sachin Shah	**Question 1:** The applicability of guidelines to all FIs in a proportionate manner is very broad and vague. It is proposed that MAS may consider making the guidance clearer by linking the proportionality to FI's risk appetite or nature of business or any other criteria deemed fit by the regulator. This will assist the FI's by providing a reference and narrow down the decision-making process for the adoption of guidelines. For example, the FI may have presence in some jurisdiction, where nature of business is not complex or jurisdictions where the risk appetite is very low due to the stringent and local regulatory obligations. **Question 2:** No comments. **Question 3:**

		No comments. The guidance provided on the proposed responsibilities of the Board and senior management in section 2.5 and 2.6 appears to be holistic in nature. **Question 4:** MAS proposal for the FIs to establish a dedicated cross-functional committee to oversee AI risk if the overall AI risk exposure of an FI is deemed material is well thought, however it may prove to be more onerous from FIs perspective as it will add to the already existing set of committees in the governance framework. It is proposed that any new guidelines should aim to assist the FIs in maintaining their overall governance framework leaner. Additionally, MAS may consider providing guidance on the addition of a subject matter expert / AI expert to the existing risk committee. This will not only assist the FIs to strike a balance between the new regulatory requirements and the existing governance framework but also help MAS in executing their supervisory obligations without any added complexities. **Question 5:** The MAS proposal for FIs to establish clear definitions, criteria and processes, supported by robust systems, to facilitate the consistent identification of AI across all relevant business and functional areas is well intended. However, MAS may consider providing guidance that these definitions/criteria/processes should not be an outcome of Generative AI. MAS may consider various instances where Generative AI is been used to copy paste the information which makes the definitions or processes more deceptive. Also, MAS may consider providing the guidance to FIs for outsourcing these tasks to consultancy industry. Let the ownership of understanding the AI ecosystem and reflecting them in the policy framework lie with FI only **Question 6:** MAS proposal for FIs to establish and maintain an accurate and up-to-date inventory of all AI usage is well intended. However, it is proposed that MAS may also consider providing guidance on the periodic review of this inventory and specify the frequency of such inventory reviews. **Question 7:** It is proposed that MAS may also consider the risk dimension of “Probability” which is related to the likeliness of the occurrence of the risk event. This is in alignment with the global standard practice of risk assessment methodologies. **Question 8:** No comments. **Question 9:** No comments. **Question 10:** No comments.
63	Silent Eight Pte. Ltd.	**Question 1:** **Question 2:**

		Question 3: **Question 4:** **Question 5:** **Question 6:** **Question 7:** **Question 8:** We recognize transparency as a critical enabler for effective AI risk management, as it allows financial institutions to identify, assess, and manage key risks arising from the use of third-party AI systems, including those related to data, model behavior, technology, and cybersecurity across the AI lifecycle. From the perspective of third-party providers, transparency can be effectively enabled through comprehensive and structured documentation covering model development, validation, performance and benchmarking, fairness and bias assessments, as well as controls for security, governance, and ongoing monitoring. Such artifacts can support financial institutions in understanding, validating, and managing AI-related risks during both deployment and ongoing use. Where deeper explainability may be constrained by the nature of certain AI techniques or the need to protect proprietary intellectual property, we support that outcome-based testing, enhanced human oversight, and appropriate user disclosures represent proportionate and effective compensatory measures to support responsible and well-controlled use of third-party AI systems. **Question 9:** We support the objective that financial institutions maintain sufficient documentation to enable reproducibility, auditability, and effective oversight of AI systems, including those sourced from third-party providers. Such documentation is important to support independent review, risk assessment, and regulatory assurance across the AI lifecycle. For third-party AI systems, we suggest that reproducibility and auditability be interpreted in a proportionate and practicable manner, focusing on the ability to reproduce and assess model outcomes, behavior, and performance under defined conditions, rather than requiring replication of the full internal development process or model construction. In this context, comprehensive documentation on data characteristics and controls, functional design and intended use, evaluation methodologies and results, explainability and fairness assessments, key assumptions and limitations, and change-management practices can provide sufficient assurance to financial institutions and reviewers. This approach would enable independent parties to understand, assess, and validate the operational and risk profiles of third-party AI systems, while recognizing differences

		between internally developed and externally provided AI solutions and offering protection of proprietary intellectual property for third-party providers. **Question 10:**
64	Singapore FinTech Association	**Question 1:** Members welcomed the approach of applying the Guidelines in a risk-based and proportionate manner according to the nature, scale and complexity of a given FI's business activities (including the manner in which the AI is used and its risk profile), and noted the importance of striking the right balance between innovation and control. In particular, members noted that AI use cases that serve assistive functions—such as copilots for writing, coding assistants, internal document summarization tools, or embedded systems for operational efficiency, which augment human capabilities rather than replace human judgment and are subject to human review and approval before any consequential action is taken—present different risk profiles and might require only lighter lifecycle controls, as opposed to AI systems that make autonomous decisions affecting customers or business outcomes or which present impersonation, misinformation, and customer confusion risks. In particular, members respectfully suggested that AI systems which do not affect customer or business outcomes (e.g. AI used in ultra-low-latency algorithmic proprietary trading) should be treated as "lower risk" in keeping with the principle of proportionate application. Members supported the general principle that the Guidelines should not apply (or apply only to a certain extent taking a risk-based approach) to "low risk" AI tools or AI use cases that are not an integrated part of business processes. This might be done by differentiating obligations by "risk tiers" (e.g. the EU AI Act's "minimal", "limited", "high" and "unacceptable" ratings). Members also respectfully noted that other regulations applying to FIs, such as the Guidelines on Risk Management Practices – Technology Risk as well as relevant Notices for different FIs (to the extent applicable), exist to address general technology risks. Accordingly, compliance with the Guidelines should not necessarily result in the creation of parallel systems, policies and procedures within the FI which might lead to inefficiency. Members also respectfully noted that specific controls, processes, and assurance mechanisms used to operationalize these principles should remain adaptive and not overly prescriptive to enable implementation approaches to evolve alongside technological developments. In view of the above and to support consistent implementation across the industry, members suggested supplementing the Annex to the Guidelines with: 1. Baseline expectations applicable across the board for all FIs; 2. Illustrative examples (particularly borderline cases such as heuristic decision engines or analytics tools with embedded AI features) of what would be viewed as "higher risk" or "higher complexity", whether according to AI use cases (e.g. those which determine customer or business outcomes), or according to the degree of integration with

		the FI's financial services or dependency of core processes on AI; and what would be viewed as "lower risk" (e.g. non-consumer / retail-facing activities); and 3. Expectations on FIs' roles and responsibilities in respect of FIs' uses of third-party embedded AI. Members also suggested short self-assessment checklists or templates to reduce potential subjective interpretation of certain aspects of the Guidelines. **Question 2:** Members made reference to their responses to question 1; namely, the request to distinguish between "higher risk" and "lower risk", models, systems or use cases, and the request to provide illustrative examples of what would and would not be an AI "use case", "system", and "model" respectively. In particular, members respectfully requested for further guidance on what would constitute a system operating on solely "predefined programming logic or rules" and for more concrete examples and edge cases on the same. Separately, members respectfully requested further guidance on whether a set of "predefined programming logic or rules" which itself was developed or modified based on AI-based learning or AI-driven data analytics, would be excluded from the scope of "AI" for the purposes of the Guidelines. Members also raised the importance of continued engagement with industry and updating of the Annex as AI grows in sophistication and evolves. **Question 3:** Members were supportive of the view of the key roles played by institutions' respective boards and senior management teams ("BSMs"). Members respectfully requested for flexibility in governance structures for AI oversight (e.g. without requiring establishment of dedicated cross-functional committees) where effective governance can already be achieved through compliance by BSMs with existing guidelines, such as the Guidelines on Risk Management Practices (e.g. specific guidelines on Technology Risk Management), Guidelines on Outsourcing and Guidelines on Individual Accountability and Conduct. Members respectfully noted that requiring cross-functional committees to be implemented even in such cases might not be proportionate. Members expressed support for ensuring BSMs receive regular training and awareness sessions to ensure, periodic capability-building for Boards (e.g. annual briefings on emerging AI risks or short training modules), not only on the strategic uses of AI, but also on malicious threats, misuse scenarios, impersonation risks, and the potential risks of AI and mitigation of these risks and challenges. To facilitate BSMs' responsibilities for implementation of AI risk management frameworks, resource allocation, and effective reporting, members respectfully suggested exploring structured reporting cadences, such as periodic AI risk dashboards and incident updates. Such periodic reporting could include not just AI system inventories and risk assessments,

		but also decision-level auditability metrics, demonstrating that AI decisions can be reconstructed and explained when required. It was respectfully suggested that in the case of FIs that are large multi-national corporations, group- or enterprise-level AI governance bodies could perform operational oversight with the BSM of the FI providing strategic oversight and challenge and permit risk-based AI literacy at a "principles-level", not providing approval at a "model-level". **Question 4:** Members were generally supportive of the idea that cross-functional committees could assist to address high AI risk materiality, and also expressed support for MAS' identified risk dimensions of "impact", "complexity" and "reliance" in paragraph 3.10 of the Guidelines. Members noted that such committees could assist to ensure proper collaboration within the organisation, a holistic understanding and use of AI in the FI, timely resolution of challenges, and escalation of risks (e.g., mis-scoring, model drift, bias events, any outages, cyber events exploiting AI components and impact on financial, operational, regulatory or conduct risk), as well as to ensure decisions are not taken in silos and that all functions of the FI understand their responsibilities related to AI governance. However, members respectfully noted the importance for flexibility of FIs to determine internal governance arrangements (including whether or not to set up a dedicated cross-functional committee) in view of other applicable regulatory guidelines such as the Guidelines on Risk Management Practices. Members noted that existing risk / technology governance committees might already cover AI-related issues (see response to Question 3 above). Accordingly, members respectfully suggested that the Guidelines express flexibility for FIs to be able to rely on existing committees (without an additional or "dedicated" committee being set up), where the roles and responsibilities of such committees cover matters which relate to AI risk (e.g. as may be set out in their terms of reference). In addition, members respectfully requested that it be considered whether, group- or enterprise-level committees might be sufficient to provide such oversight, without a committee of the Singapore FI specifically being additionally established (provided BSMs of the Singapore FI remain represented at such group- or enterprise-level committees). Separately, members respectfully suggested that in determining the materiality of AI risk, the Guidelines could provide further guidance on the parameters for assessing "overall AI risk exposure" (ideally aligned to international norms in this space, e.g. the EU AI Act). FIs could rely on a common set of artefacts such as an AI Inventory, a "materiality register", stage-gate evidence, and outcome review documentation, and could also rely on a set of indicators such as: • percentage of critical processes reliant on AI; • number of high-risk AI use cases; • concentration of AI dependencies on external vendors; • volume of customer interactions influenced by AI.

		Members also noted indicators of "risk materiality" might differ for FIs that do not face consumers or retail customers, and might include factors such as (a) the effect on market integrity; (b) implications for operational resilience of the FI; and (c) risks of model failure, rather than any detriment to "consumers". Members also respectfully requested for MAS to provide any further guidance (e.g. guidelines or best practices) on how it would expect such committees to be composed and the scope of their decision-making authority. Members noted that the efficacy of cross-functional committee might vary depending on various factors, including but not limited to access to comprehensive AI risk information and the strength of their mandate and authority. **Question 5:** Members were generally supportive of MAS' proposal for FIs to establish consistent AI definitions, criteria, processes and systems as a part of an FI's broader responsibility to build a culture prioritising security, governance and data protection, and acknowledged the importance of periodic review / ratification / attestation of such definitions, criteria, processes and systems by the FI's BSM. However, members observed that existing requirements applicable to FIs might suffice to address these matters. For example, in the Guidelines on Risk Management Practices – Technology Risk, members noted requirements on management of information assets which requires FIs to identify and inventories information assets, which could include AI models (see response to Question 3 above). On this basis, members respectfully requested that MAS consider whether the Guidelines could provide for the flexibility of FIs to leverage and enhance their existing databases or inventories of applications (whether at the entity-level or at the broader global, group- or enterprise-level), instead of establishing new systems which might overlap with those already existing. Members also respectfully requested that the Guidelines distinguish between identifying such definitions, criteria, processes and systems in relation to production-grade AI models (e.g. used in live trading, monitoring or risk controls), and less material research artefacts (e.g. experimental models, prototypes, or back-testing experiments). Members respectfully noted that requiring identification of the latter (i.e. less material research artefacts, which might be far more voluminous), would lead to a substantial and costly administrative burden without necessarily improving AI risk management. Member respectfully suggested the Guidelines set out a sample identification workflow, which could include the use of central registration processes and tooling (e.g. fields in project initiation forms, model risk systems, or IT/service catalogues) for a centralised inventory. Members also respectfully suggested a taxonomy to assist in such AI identification. To elaborate, in addition to the distinctions between "models", "systems" and "use cases", distinctions can be made between "internal" and "third-party" AI, as well as "generative" and "predictive" models. **Question 6:** Members were generally supportive of MAS' proposal for FIs to establish AI inventories as an operational control and subjecting such inventories to periodic presentation and review

		by the BSM (with updates to account for new AI systems or significant modifications), but respectfully reiterated their response to Question 5; namely that existing requirements applicable to FIs might suffice to address these matters (see above). On this basis, members respectfully requested that MAS consider whether the Guidelines could provide for the flexibility of FIs to leverage and enhance their existing databases or inventories of applications (whether at the entity-level or at the broader global, group- or enterprise-level), instead of establishing new systems which might overlap with those already existing. Members also respectfully requested that the Guidelines distinguish between deployed production-grade models (which should be included in such an AI inventory) versus less material research artefacts (e.g. exploratory research), and that the inventories be scalable to small teams. Members also respectfully proposed that the Guidelines allow the flexibility for FIs to define their own relevant materiality thresholds for inclusion in their respective AI inventories, and where possible for some guidance to be provided as to the inventory schema / fields (e.g. owner, risk classification, model type, data sensitivity, vendor details, validation status, and monitoring results) expected to be included. Under such a framework, members' AI inventories would include largely those which are assessed to be higher risk (rather than functioning as an exhaustive list of all AI usages). Members also noted the importance of reviewing and updating such inventories. Apart from controls at the model layer, members also noted the importance of controls at the data layer to independently validate results, obtain internal assurance, and exercise continued oversight on AI usage. Particularly where AI-generated structured data flows into downstream processes. Members noted the importance of trust or provenance markers to identify AI risk, retain its context, and prevent "invisible propagation". **Question 7:** Members were generally supportive of the risk dimensions identified, and noted that other dimensions might include: • Third-party concentration / dependency / jurisdiction risk (e.g., reliance on a small number of AI providers / APIs or foundation models, and/or where such third-parties are concentrated in a small number of foreign jurisdictions); • Security amplification (e.g., prompt injection, misuse of tools, unintended data exposure); • Human factor considerations (e.g., oversight fatigue or deskilling); • Explainability (e.g. ability to clarify how a model makes decisions in a way that is accessible to end users – see the European Data Protection Supervisor's Guidance for Risk Management of Artificial Intelligence Systems: https://www.edps.europa.eu/system/files/2025-11/2025-11-11_ai_risks_management_guidance_en.pdf) ("EDPS Guidance"); • Data integrity / security / sensitivity (especially where highly sensitive or regulated data is used); and • Impact on downstream systems / output. **Question 8:**

		Members were generally supportive of the proposed standards, processes and controls in paragraph 4 of the Guidelines, and their application in a proportionate manner which scales with materiality and use within the organisation. This allows flexibility for FIs to determine the appropriate standard, process and control (as the case may be), depending on the risks presented (e.g. frequency of monitoring cadences). While members noted that life cycle phases (e.g. design, build, test/validate, deploy, monitor, retire) and junctures for approval and monitoring should be defined clearly, members noted the importance of adopting a principles- and outcomes-based approach rather than prescribing specific standards, processes and controls for AI life cycle management. Members did note that many AI-driven workflows might not preserve provenance and validation metadata at the data level and accordingly may affect FIs to trace a specific decision or reportable figure back to its originating evidence, model version, and validation state. Members considered whether explainability could be more robustly framed to include not only model-level explainability but also "decision-level" explainability (e.g. storing prompts/of prompt hashes, file hashes, model seeds, model versions used, complete output hashes, pages number/original values where data was read from (in the case of document processing), what business rules were applied to which field, confidence values, what data was manually verified or reconciled, etc.). However, members also noted that explainability requirements vary significantly: systems with low risk may not require explainability mechanisms, while AI systems with output that may be used in a manner that could impact safety will likely need a method for explaining or providing insights into how the system performed its analysis or the factors influencing outputs. Other areas in which the principle of "proportionality" applies, in determining the degree of standards, processes and controls to be implemented, include but are not limited to determination of (a) human oversight (not only of outputs but also of underlying evidence, confidence and validation history); (b) testing (e.g. functional, adversarial, compliance testing); and (c) third-party risk. For example, AI life cycle controls might differ between first-party AI use cases, systems or models developed and deployed by the FI itself, versus FIs that deploy AI developed by third-parties. In relation to the latter, members respectfully sought further guidance from the Authority on how far FIs should be held responsible for aspects of the foundational AI service controlled by third-parties, and suggested that it is the third-party AI provider (rather than the FI) that should remain responsible for matters such as model safety and security, training data quality, model architecture, infrastructure resilience, platform security, vulnerability management, and compliance with applicable technical standards. In specific relation to paragraph 4.4 of the Guidelines, members respectfully sought further clarity on the expectations surrounding "contingency plans" for AI assessed as high risk, noting that AI system reliability might be addressed in multiple ways, for example by using alternative models, reverting to backup processes during model failure, or ensuring quick time-to-recovery during platform outages.

		In specific relation to paragraph 4.19 of the Guidelines, members respectfully sought further clarity on (and examples of) the extent of "independence" expects of the parties undertaking assessment and review prior to deployment of AI use cases, systems or models, apart from them needing to be "independent from the development and deployment teams [of FIs]". More generally, members also respectfully noted that some of these standards, processes and controls might already be implemented by FIs as part of their compliance with existing MAS Notices and Guidelines. For example, at the model layer, some companies might already follow rigorous model risk management practices in compliance with existing regulatory requirements and/or guidelines. This may include already maintaining a comprehensive governance framework of enterprise policies and controls, for model development, documentation, validation, usage and ongoing monitoring plans. Notwithstanding this, members did respectfully request for further illustrative practical examples of (a) testing standards for high-impact models; (b) transparency disclosures for customer-facing AI; (c) recommended monitoring KPIs (e.g. control coverage rates, incident response times, risk trends for reporting); and (d) expectations for kill-switches or fallback mechanisms, potentially by way of templates or reference materials available through Project MindForge. **Question 9:** Members raised other forms of risk management not covered in the proposed Guidelines, including — 1. Obtaining attestations from AI providers on training data provenance, model assurance standards, toolchain hygiene, to strengthen supply chain integrity; 2. Establishing policies to govern AI orchestration to preserve control from Singapore when coordinating AI capabilities across jurisdictions; 3. Ensuring employee readiness; 4. Considering "interoperability" with international frameworks, i.e. compliance with internationally recognized AI standards and frameworks and the degree to which such compliance is equivalent to compliance with the Guidelines (e.g. the EU AI Act, ISO/IEC 42001 Artificial Intelligence Management System certification); 5. Implementing controls to prevent AI misuse (e.g. access, operation and supervision of AI systems within organisations, internal verification, authentication, and multi-level authorization); 6. Keeping books and records of AI implementation and adoption, including clear rationale(s) for why AI adoption is required (i.e. an input-output-desired objective analysis); 7. Capacity building / skills development / enhancing expectations around continuous education and user awareness (i.e. organisational training on risks of AI misuse, over-reliance, and misinformation). Members respectfully noted the importance of distinguishing between applications of generative AI on the one hand, and other statistical- or rules-based programmes on the other (e.g. those generally responsible for core transaction execution and decisioning functions), which present a different risk profile. Members respectfully sought more clarity

		on whether statistical- or rules-based programmes would fall within the Guidelines' conception of "AI". Relatedly, members respectfully requested the applicability of the proposed Guidelines to focus more on higher-risk generative AI use cases, as well as for the Guidelines to incorporate flexibility (whether through a "removal of difficulties" clause or otherwise) for FIs to address case-specific difficulties and challenges with implementing the Guidelines other than through what might expressly be prescribed in the Guidelines. Separately, members also respectfully requested that MAS engage with the Personal Data Protection Commission to consider whether the proposed Guidelines would be consistent with the obligations under the Personal Data Protection Act 2012 (e.g. regarding data residency, data minimization, or the "right to explanation" for automated decisions). Members suggested that in relation to human oversight and third-party AI governance, such matters should be implemented on a risk-based approach with respect to the potential harm to be caused to the customers, rather than for all AI use cases. **Question 10:** Members generally requested for a consider a longer period for implementation (up to 18 to 24 months), in view of international developments and the need for larger organisations (particularly those that have global operations) to align on their systems, policies and procedures across different jurisdictions. Members were also generally supportive of a phased implementation approach, suggesting such "phases" to allow for higher-risk or more "material" AI applications to be addressed first. In general, such a "phased approach" could involve: 1. Establishing the AI inventory and "materiality" register; 2. Implementing tools / rule mapping and standards, processes and controls to address AI risk; and 3. Operationalising compliance with Guidelines by rolling out policies and procedures and conducting employee readiness / training.
65	Singapore Reinsurers' Association	**Question 1:** We support MAS' proposal to apply the Guidelines to all FIs, together with the emphasis on proportionate implementation. We agree that application of the Guidelines should be commensurate with the size and nature of the FI's activities, its use of AI and its risk profile, as set out in para 1.5. We note that the Annex sets out guiding questions to determine whether AI is used as an integrated part of the FI's business processes, including whether the • lack of access to AI services or tools would disrupt workflows that the FI is materially dependent on, or • if AI is integrated with systems that the FI is materially dependent on. We suggest that these guiding questions be applied as an initial determination of applicability. Thereafter, consistent with para 1.5, the expectations in Sections 2 and 3 of the Guidelines should continue to be applied based on the relevance and risk materiality of the specific AI use case, system or model. Reason being, FI's material dependency on AI

		for business processes does not, on its own, imply that the AI use case presents material AI risk. This risk-proportionate approach enables flexibility while ensuring that governance and controls are appropriately calibrated. **Question 2:** We support the proposed scope of the Guidelines, including the definition of AI by reference to AI models, systems and use cases and the focus on models or systems that learn and/or infer from inputs to generate outputs. We welcome MAS' clarification in para 1.2(d) that 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. We agree with MAS in para 1.4 that AI risks vary depending on the specific AI use case, and that the application of the Guidelines should therefore take into account the AI use case rather than be applied uniformly. **Question 3:** We support the proposed responsibilities of the Board and senior management in overseeing AI risk management where AI is used as an integrated part of business processes. **Question 4:** With regards to the proposal to establish a dedicated cross-functional committee, our view is that this should be left to the discretion of the Board and FI on what governance structure to put in place that would support their overall obligation to manage AI risks and discharge oversight responsibilities. **Question 5:** We support MAS' proposal for FIs to establish clear definitions, criteria and processes, to consistently identify AI usage across relevant business and functional areas. We also support the expectation in para 3.3 for FIs to assign clear roles and responsibilities for AI identification. We agree that processes should be regularly reviewed and updated to take into account newer AI technologies, and implemented in a manner proportionate to the nature, scale and risk profile of the FI's AI use. **Question 6:** We support the proposal for FIs to establish and maintain an accurate and up-to-date inventory of all AI use cases. We welcome the flexibility for branches and subsidiaries to leverage group-level AI inventory. This flexibility supports implementation approaches that are aligned with our existing systems and processes, while ensuring relevant AI use cases, systems and models are captured and maintained.

		We also agree that the scope and level of detail of the AI inventory should support effective oversight and be commensurate with the relevance and risk materiality of the AI use cases, systems or models. **Question 7:** We support the use of Impact and Reliance as key risk dimension for assessing the materiality of AI risks. However, while we agree that considerations relating to complexity can be relevant in informing the design and intensity of risk mitigants (e.g. explainability), we suggest treating complexity as a contextual factor rather than a primary determinant of AI risk materiality. Three reasons: 1) Complexity does not drive risk outcomes on its own. Highly complex models can be low risk if used for benign purposes (e.g. summarisation), while simpler models can be high risk if used in critical decisions. 2) Complexity evolves very quickly. What is "complex" today can become routine tomorrow. 3) Complexity is inherently subjective. The same model may be "complex" for one FI and manageable for another, depending on expertise, tooling and operating maturity. We further note that AI risk materiality assessments are most effective when conducted in a holistic, use-case-based manner, involving relevant functions as appropriate. This enables risks to be assessed in the context of how AI is used. Where material changes occur to an AI use case or system, a reassessment should be triggered, with the scope of reassessment calibrated to the nature and impact of the change. Overall, we consider that a risk materiality assessment framework that prioritises outcome-based dimensions (impact and reliance), supported by contextual considerations such as complexity, would better support practical, proportionate and consistent application of the Guidelines across the industry. **Question 8:** Third-party AI risk management With respect to the management of third-party AI risks in para 4.11, we support MAS' objective of ensuring that FIs have adequate onboarding, development and deployment controls for third-party AI. We agree that risks arising from the use of third-party AI should be identified, assessed and managed appropriately. In relation to para 4.11(a), we consider that effective assessment of transparency from third-party AI providers would be better supported by more consistent and standardised disclosures by such providers. This would help establish clearer and more consistent transparency expectations across FIs. It would also reduce fragmentation arising from differing information requests by individual FIs. In this regard, we note that international regulations, such as the EU AI Act's transparency obligations for general-purpose AI (GPAI) providers, may offer helpful reference points when considering how such disclosures could be structured at an ecosystem level.

		In relation to para 4.11(c), we note that assessing where the data used to train an AI model came from, integrity and related supply chain risks for third-party AI models may be subject to practical limitations. Information on model development processes and training data is often proprietary and not fully disclosed to FIs. In such circumstances, we consider it important that FIs are able to rely on reasonable due diligence measures, contractual arrangements and compensating controls, as provided for in para 4.11. Accountability for risks should be assessed in a manner that is proportionate to the information reasonably available to the FI. Overall, we support MAS' life-cycle-oriented and risk-based framework in Section 4 of the Guidelines. We emphasise that the relevance, depth and form of the proposed controls, including technology and cybersecurity reviews, should continue to be determined by the risk materiality of the specific AI model, system or use case, to ensure practical and proportionate implementation. **Question 9:** In applying the Guidelines, we encourage MAS to ensure that the proportionality principles in paras 1.5 and 1.6 are consistently reflected in supervisory practice. In particular, while AI use that is materially embedded in business processes, material dependency should not be taken as equivalent to material AI risk. Supervisory expectations should continue to be determined by the relevance and risk materiality of the specific AI use case. In relation to Section 5, we encourage MAS to apply expectations on AI training and capacity-building in a manner proportionate to the materiality of AI use and the FI's role, and to treat these as risk-based good practices rather than uniform requirements. **Question 10:** We support MAS' proposal to provide a transition period for the implementation of the Guidelines. Given the breadth and depth of the proposed expectations, we suggest that MAS consider allowing more time to allow for phased-in implementation, to support effective and orderly implementation.
66	Skandinaviska Enskilda Banken AB (publ) Singapore Branch	**Question 1:** **Question 2:** We request for MAS to consider whether the AI model and AI systems definitions can be enhanced or clarified. At present, "AI model" is defined too generically and may also be applicable to a traditional model that has no AI components in it. Similarly, "AI systems" definition at present does not clarify the technical components or machine-based components that make a system an "AI" system and that can clearly be differentiated from a traditional/normal system. **Question 3:** **Question 4:** **Question 5:** The MAS definitions of AI model and AI system differ from the EU AI Act. Similarly, the MAS definition of whether the AI would be regarded as being an integrated part of the

		business processes in the bank and inside scope of the proposed Guidelines differ from the EU AI Act Risk categories (Prohibited, High Risk, Transparency and Minimal Risk). Given the different definitions, this would create complexity and difficulty for European incorporated banks in Singapore when it comes to consistent identification/categorization of AI in the bank, maintain up-to-date inventory of AI usage, and conduct risk materiality assessments. Could MAS consider exploring greater alignment with other global regulatory AI frameworks and definitions? **Question 6:** **Question 7:** **Question 8:** **Question 9:** **Question 10:** Given the complexity involved and the challenge of alignment across regulatory requirements in multiple jurisdictions, could MAS consider a longer transition period of 24 months?
67	SonarSource Sàrl	**Question 1:** In the context of AI-assisted software development, we advise refining the definition of "proportionality." The "Source-Agnostic" reality of modern coding means that code issues (reliability or maintainability), vulnerabilities and 3rd party software libraries in low-risk tools can often be pivoted into a high-risk exploit. E.g. CWE-401: Missing Release of Memory after Effective Lifetime. Most memory leaks result in general product reliability problems, but if an attacker can intentionally trigger a memory leak, the attacker might be able to launch a denial of service attack (by crashing or hanging the program) or take advantage of other unexpected program behavior resulting from a low memory condition. Therefore, we propose an absolute standard of verification across all criticality levels: every line of code, whether written by a human or an AI, must undergo the same baseline automated code verification analysis. Proportionality should strictly apply to the consequence of findings (e.g., blocking deployment for mission-critical systems vs. logging warnings for innovation sandboxes), rather than the depth of the scan itself. References: https://cwe.mitre.org/data/definitions/401.html **Question 2:** We recommend the Guidelines explicitly distinguish between "Assistants" and "Agents". In software development, AI Agents have access to Information Assets and are increasingly capable of writing, testing, and merging code autonomously, while Assistants offer suggestions or code snippets, requiring the human developer to actively review, accept, and integrate the output.

		This shifts the risk profile for AI Agents from bad advice to malicious action. We advise that any AI Agent with write-access to internal code repositories be automatically classified as High Materiality, requiring a mandatory, non-negotiable "Human-in-the-Loop" protocol based on a pre-defined set of significant, risk-specific checkpoints to validate business logic. Also, Agents must not have unrestricted access. FIs must enforce strict permission boundaries, ensuring agents function with the minimum privileges necessary (e.g., restricting access to specific repositories and disabling external internet access if not strictly required). This ensures the force multiplier effect of AI does not become a risk multiplier. We also recommend logging facilities to be mandatory to ensure agent actions are recorded, verified, non-disputable and non-repudiable. **Question 3:** We suggest that the senior management's responsibility for resource allocation be highlighted as the primary enabler of compliance. Adopting AI coding tools without funding the corresponding defensive infrastructure creates an accountability gap. Effective oversight requires investment in centralized verification systems to support verification processes that can scale. Proof of Feasibility: We point to the example of one of the world's largest global banks, which successfully supports thousands of developers by utilizing a centralized and independent code verification solution. By funding this infrastructure, they verify 100% of new code with homogenous and regulatory aligned assessments across the enterprise. Without similar Board-level commitment to infrastructure resources, FIs will be unable to maintain safe AI adoption at scale. **Question 4:** In the context of AI-assisted or agentic software development, we recognise the level of integration of these assisted or agentic systems into an FI's Information Assets and operations. Therefore, we submit that for this application areaAI risk should effectively be considered material for all modern FIs. Applying the Source-Agnostic, Risk-Specific framework, we advise that the Committee's mandate should focus on distinguishing between verifiable and subjective risk rather than simply assessing exposure. The Committee must categorize the organization's AI usage based on the nature of the output. For verifiable outputs, such as software code, the result is deterministic and can be rigorously proven correct via automated testing; here, the Committee must mandate strict verification of output protocols. In contrast, for subjective outputs like natural language advice, where rigorous automated testing is impossible, the Committee must explicitly define risk appetite and mandate human-heavy oversight. By owning this distinction, the Committee ensures that rigorous engineering controls are applied where they are most effective, while allocating scarce human oversight resources to the areas that truly need them. **Question 5:**

		We advise FIs to refine their criteria by explicitly distinguishing between High-testability domains and low-testability domains. Low-testability domains: In domains where outputs are subjective or social (e.g., a chatbot potentially giving biased investment advice or using discriminatory language), rigorous automated testing is difficult. Here, consistent identification is the primary control. You must know exactly where these models operate to apply human oversight and monitoring. High-testability domains: In engineering, the output (code) is highly verifiable. We can prove if code is syntactically correct, secure, and functional. In this space, the criteria should not focus on identifying usage (which is effectively detecting the tool), but rather on verifying outputs. For software development, the question shouldn't be "Did you use AI?" but "What testing are you doing of the output?" By focusing criteria on the verification process (SAST, unit testing) rather than just identification, FIs leverage the deterministic nature of code to ensure safety, rather than getting bogged down in administrative tracking of AI usage. **Question 6:** We advise FIs to approach the inventory requirement with an already mentioned source-agnostic mindset. While we support the principle of transparency, maintaining a static inventory of tools is operationally difficult given the speed of AI evolution. To future-proof the institution, we must govern the artifact (the code), not the origin (the tool). The standard for code quality and security remains absolute. It is irrelevant whether a line of code was written by a human, an established LLM, or a brand-new experimental agent. All code must meet the same rigorous engineering standards before entering the production environment. The operational principle is straightforward: where rigorous testing can be applied, be source-agnostic. Rather than expending resources on tracking usage, FIs should focus entirely on testing the output. This prevents wasting time on less important metrics and ensures the approach remains future-proof, regardless of how many new tools, sources, and formats of interactions with AI appear on a daily basis. **Question 7:** We advise refining the definitions of complexity and reliance to reflect the specific risks of Generative AI in coding: Complexity: Data shows that 61% of developers agree AI produces code that "looks correct but isn't reliable." Materiality assessments must account for this deceptive quality, where code appears functional but contains subtle logic flaws. Reliance: Despite high adoption, 96% of developers do not trust AI-generated code. However, time pressures often force them to accept it. Materiality assessments should capture the risk of automation bias overriding developer judgment

		Security consciousness: We strongly urge MAS to recognize that AI models demonstrate a fundamental lack of security consciousness. This is not a matter of occasional hallucination, but a structural failure rooted in the foundational design.
		
		>>> LLMs struggle to prevent critical issues like injection flaws because doing so requires taint-tracking (analyzing data flow from an untrusted source to a sensitive sink) - a non-local analysis often beyond their context window. Furthermore, because they predict tokens based on training data, they frequently reproduce "hardcoded secrets" (passwords, API keys) present in that data.
		
		>>> Our analysis indicates that for non-reasoning models, a majority of generated vulnerabilities are of 'BLOCKER' severity (the highest possible rating). Therefore, materiality assessments must assume that all raw AI output is inherently insecure by design until verified.
		
		Determinism: We suggest adding "Determinism." Unlike standard software, AI coding tools are non-deterministic (generating different code for the same prompt). Assessing whether a system's behavior can be reliably reproduced is a critical risk dimension for financial software compliance.
		
		References: https://www.sonarsource.com/the-coding-personalities-of-leading-llms.pdf
		
		Question 8:
		In the context of software development, we observe varying degrees of code analysis maturity across FIs globally. To assess maturity, we reference a 5-level framework:
		
		Level 1: External audits - FIs rely on external specialists to review code
		Level 2: Manual review & spot checks amongst the team - Individual developers, or teams verifying code without a mandate from leadership - usually without a common standard
		Level 3: Automated code checks - Code is reviewed by a software system during PR or build
		Level 4: Standards enforcement - A common standard for software development is established
		Level 5: Standards governance - Common standard is well governed across teams and projects
		
		The historical disparity in software quality & security standards across firms is being further exacerbated by the high-frequency output of AI agents, which creates a volume of code that outpaces traditional controls that could evolve to new risk surfaces. Considerations & recommendation as follows:
		
		>>> Manual Review: While essential for logic, manual review cannot scale. With the effort required to review AI code being significantly higher (38% more effort, per recent industry data), relying on humans as the primary gate is mathematically unfeasible and error-prone.
		
		>>> Dynamic Testing/DAST: While accurate, DAST occurs too late in the lifecycle. Remediation at this stage is costly and slows velocity.

		>>> Automated Code Checks (Static Analysis/SAST of known bad patterns): We propose that FIs need to achieve Level 3 code analysis maturity upon implementing AI coding assistants and automated Static Analysis be recognized as the required primary control. It is the only mechanism that scales linearly with AI generation. >>>>> Proof of Impact: Data shows that teams utilizing an independent automated static analysis verification layer are 24% more likely to reduce vulnerability rates and 20% more likely to reduce defect rates. Furthermore, large enterprise financial services companies have demonstrated that this standardized analysis leads to a measurable increase in code consistency and maintainability. Aligning code inspection velocity with AI-driven development is the first step. To meet operational resilience requirements of MAS (minimise unexpected software-led outage risks or delayed patching) closing the 'remediation gap' is a crucial second step to ensure that the increased volume of defects inherent in agentic development is resolved before manifesting as service disruptions or compliance breaches. To achieve this, we propose integrating Automatic Program Repair (APR) tools to augment manual remediation and maintain a secure, stable production environment, under human oversight. Agentic coding assistants and systems (including their safety guardrails) are expected to become more performant and resilient over time. To ensure that the chosen technology remains 'fit-for-purpose' under MAS's standards, we recommend that FIs periodically reference publicly available purpose-built benchmarks for code quality and security issues. This data-driven selection process ensures that model upgrades are not merely based on speed, but on a measurable improvement in code quality and security reliability. For FIs opting to deploy open-source models for their AI coding assistants, the responsibility for maintaining performance, resilience, and 'fit-for-purpose' status - as defined by MAS, lies entirely on the institution. We submit that these models must be subject to continuous benchmarking throughout their lifecycle. Furthermore, FIs must establish robust pipelines to ensure coding assistants are updated in a timely manner, preventing the emergence of zero-day vulnerabilities or 'model drift' that could compromise the security of the generated code. The adoption of AI coding assistants by third-party providers introduces secondary supply-chain risks to MAS-regulated entities. We propose that the same rigorous standards for code quality and remediation be integrated into the Third-Party Risk Management (TPRM) framework during onboarding and annual reviews. Crucially, while vendors are expected to adhere to these benchmarks, the ultimate accountability for operational resilience remains with the FI, necessitating a 'trust but verify' approach to third-party code integrity. References: https://www.mas.gov.sg/regulation/notices/notice-fsm-n05 https://www.mas.gov.sg/regulation/notices/notice-fsm-n06 https://www.sonarsource.com/the-coding-personalities-of-leading-llms/leaderboard/ **Question 9:**

		We advise that the Guidelines address the risk of regulatory obsolescence due to the extreme speed of AI evolution. New models, approaches, and tools appear daily. A governance model that relies on defining approved sources will be obsolete before it is published. The Guidelines should explicitly endorse a "Source-Agnostic" approach, where governance focuses on the quality of the output (the code) rather than the nature of the input (the tool). This future-proofs the regulation against the rapid emergence of new AI technologies. Additionally, we suggest the Guidelines address "False Positive Rate" as a critical success factor for automated controls. Simply implementing a tool is insufficient; the tool must be precise. If an automated enforcement mechanism floods developers with incorrect alerts, it leads to alert fatigue, causing developers to ignore valid security warnings. High-precision tooling is a prerequisite for effective governance in a high-velocity environment. **Question 10:** Regarding the specific use case of AI-powered software development, we advise that a 12-month transition period presents a significant window of exposure. The risk of AI in coding is not a future event; it is a present reality. Developers across the sector are arguably already utilizing these tools, often outside of formal governance. To delay the enforcement of verification standards by 12 months is to implicitly accept that a year's worth of potentially unverified, probabilistic code may accumulate in the foundations of the financial system. We urge MAS to treat the deployment of automated verification as an immediate operational necessity rather than a medium-term target. Instead of a phased transition, we recommend a posture of heightened awareness and accountability starting immediately. FIs should be encouraged to activate their defensive safety nets (verification tools) as soon as possible to identify and mitigate the risks that are being generated today, ensuring that governance paces the reality of adoption.
68	Sumitomo Mitsui Banking Corporation	**Question 1:** N/A **Question 2:** N/A **Question 3:** N/A **Question 4:** In relation to the assessment of the materiality of an FI's overall AI risk exposure mentioned in the 2nd last sentence of paragraph 2.4 of the draft Guidelines, footnote 12 refers to paragraphs 3.8 to 3.11. However, the focus of paragraphs 3.8 to 3.11 is risk materiality of AI use cases rather than an FI's overall AI risk exposure. Could MAS provide more clarity or specific guidance in relation to determining what is considered material overall AI risk exposure?

		Could MAS clarify whether for determining what is considered material overall AI risk exposure at the organisational level, FIs should aggregate individual use case materiality or whether other organisational-level factors should also be considered? We would propose that the Guidelines set the expectation for FIs to have their materiality assessment of their overall AI risk exposure performed by a cross-functional committee or working group so that the assessment may be adequately performed with the involvement of persons possessing the relevant expertise. **Question 5:** Could MAS clarify whether the usage of support systems is mandatory for the identification of AI usage? **Question 6:** N/A **Question 7:** Could MAS clarify the definition of complexity of an AI use case? Banks have established robust and mature frameworks for evaluating the significance of business processes, data criticality, model tiering, etc. Can MAS articulate whether these should be factored in when determining the risk materiality of an AI use case? **Question 8:** Could MAS clarify whether the AI lifecycle is expected to be aligned with existing governance policies, including those governing model, data, or system lifecycles? **Question 9:** The draft guidelines are more focused on 1LoD controls than overall risk management oversight. Could MAS consider including components to AI risk management at the enterprise level? **Question 10:** We would propose a longer period of between 18 to 24 months, as 12 months or so would be required to set up policies/ frameworks/ processes. After the 12 months for setting up policies/ frameworks/ processes, further time will be required to operationalize and stabilize AI governance and oversight.
69	Tanveer Bhatti	**Question 1:** I support applying AIRG to all FIs with expectations calibrated to their use of AI, as set out in the main text and the Annex. MAS already distinguishes between AI used as an integrated part of business processes and AI used in ad hoc assistive ways and provides examples of each in the Annex. Aligning control expectations explicitly with those tiers would make proportionality more concrete. For use cases that are not integrated into business processes, such as occasional use of an LLM to draft emails under human review, basic policies are likely to be sufficient, including clear rules on allowed and disallowed uses, prohibitions on inputting confidential or client data into public tools, requirements for human checking before use, communication to staff and periodic checks for compliance.

		For AI that is integrated into business processes but does not support critical services or high impact regulated activities, it is reasonable to expect firms to produce a standardised Assurance File and operate ongoing monitoring, with independent challenge proportionate to risk. That challenge can be delivered through sampling, thematic reviews, or validation at model family level, rather than a full bespoke validation for every instance as a default. For integrated, high materiality AI such as models used in critical business services, regulatory capital, AML and fraud, or regulated advice with significant customer impact, the same core Assurance File can be combined with full independent validation and portfolio level review, consistent with MAS' expectations on high materiality models and independent re validations. Documentation and governance routing can consume significant specialist capacity for high materiality AI use cases. Standardised, banded Assurance File expectations would allow firms to redirect effort toward testing, red teaming, monitoring, and customer outcome analysis rather than bespoke document production. Efficient governance, where risk outcomes are demonstrably met, is a hallmark of sophisticated risk management rather than a relaxation of standards, and AIRG is an opportunity to make that expectation explicit. These minima are not safe harbours. FIs remain fully accountable for outcomes, including consumer harm or systemic impacts, regardless of documentation compliance. Clear minimum Assurance File expectations clarify process adequacy and improve comparability across firms and use cases, while preserving MAS' discretion to assess outcomes and take action where harms occur. **Question 2:** I support the scope MAS has set out, which covers models, systems, and use cases where models learn or infer from inputs to generate outputs such as estimates, predictions, content, summaries, recommendations or decisions, and explicitly excludes calculators or tools whose outputs are solely based on predefined programming logic or rules. Hybrid solutions that combine deterministic rules and AI components can naturally be treated as in scope for the AI elements, without reclassifying deterministic components as AI purely because they share a system boundary. A recurring source of operational risk is AI appearing silently in third party systems via upgrades or toggleable features. Treating activation of AI features in vendor products, such as AI assist modes or upgraded screening modules, as an AI identification trigger would directly address this. FIs can reasonably be expected to seek contractual provisions requiring vendors to disclose when AI is introduced or materially changed in products they use, and to treat the decision to enable such features as a governed deployment rather than an operational tweak. Where adequate disclosure or pre notification cannot be obtained because of vendor leverage or standard terms, it is prudent to apply compensating controls. Those include constrained scope of deployment, increased testing and monitoring frequency, shorter review cycles, and treating the dependency as higher materiality unless and until evidence supports a lower rating. **Question 3:**

		MAS emphasises that the Board and senior management must maintain effective oversight of AI related risks, foster an appropriate risk culture, and ensure that AI use does not conflict with other supervisory expectations, integrating AI into existing risk management frameworks, policies and practices across the organisation. That emphasis is appropriate. Oversight becomes more auditable and less dependent on organisational form if each FI designates a named senior executive accountable for AI risk, with clear responsibility for implementing AIRG and overseeing AI risk across the firm. AI risk appetite can be explicitly integrated into existing model risk, technology risk, operational risk, conduct, and outsourcing frameworks, for example through limits on automated decision rates, performance drift thresholds, dependencies on single providers, and metrics on complaint and overturn rates for AI driven decisions. Implementation of AIRG through existing model risk, technology risk, operational risk, compliance and internal audit structures, with additional AI specific capability where scale and complexity justify it, is typically more efficient than building parallel structures. Concise logs for approvals, exceptions, incidents and near misses for material AI use, including where Board or committee decisions have been made, provide the traceability supervisors need without proliferating additional meetings. **Question 4:** MAS proposes that where overall AI risk exposure is material, FIs establish a dedicated cross functional committee to ensure adequate oversight and to address potential gaps in risk management coverage. The underlying objective is cross functional control of material AI risk. That objective can be met in different ways. A dedicated AI committee is one valid implementation. In many institutions, existing enterprise level committees such as model risk, operational risk, or technology and change already operate with cross functional membership and defined authority. Where such bodies can demonstrate representation of business, risk, compliance, technology, data and operations, formal authority to approve, limit or halt material AI deployments, and a defined cadence and agenda that include the AI inventory profile, materiality distribution, incidents, and portfolio dependencies, they can meet the same supervisory intent without creating additional forums. The intention is not to weaken the expectation of heightened governance for material AI exposure, but to recognise that many FIs already operate cross functional committees with equivalent remit and authority. Where those existing bodies can demonstrate that they perform the same functions MAS envisages for a dedicated AI committee, the supervisory outcome is functionally identical, without creating duplicative forums. Assessing organisational level AI exposure benefits from explicit portfolio indicators rather than only counting high materiality use cases. Relevant indicators include the concentration of critical AI functions on particular models or providers, the dependency of critical business services on those providers, and exposure to synchronised updates or outages across multiple use cases. High level clarity on how the industry is expected to respond when critical vulnerabilities in widely deployed models are identified, including when temporary restriction, heightened monitoring, alternative controls or kill switch activation at portfolio level are expected, would further strengthen systemic readiness. **Question 5:**

		I support MAS' expectation that FIs establish systems, policies and processes to ensure consistent identification of AI usage across all relevant business and functional areas. Identification becomes reliable when it is driven by clear triggers. Vendor introduced AI is a natural trigger, so that any notification from a vendor that AI has been added, enabled, or materially changed in a product used by the FI automatically initiates AI identification, inventorying, and materiality assessment. It is also helpful to state explicitly that AI used in regulated activities such as credit decisioning, insurance underwriting, investment advice and portfolio management, or in AML, sanctions and fraud controls, and AI with access to sensitive customer or employee data, is always within scope for identification, inventory and materiality assessment, even if described internally as assistive or pilot functionality. Assigning a control function with authority to require declarations of AI use and arbitrate borderline cases significantly reduces the risk of shadow AI and inconsistent application of AIRG. Passive identification based only on declarations and vendor notifications is not sufficient in high-speed environments. In practice, AI often enters via unannounced feature additions in standard software or individual staff using unapproved browser based tools. A credible identification regime therefore also relies on active discovery, proportionate, risk based, and consistent with applicable privacy and monitoring policies. For example, periodic reviews of network traffic for calls to known AI providers, scanning for unapproved AI related domains or APIs, and shadow IT audits focused on AI enabled tools. This closes the gap between stated policies and actual usage and reduces the likelihood that material AI operates entirely outside formal governance. **Question 6:** I support MAS' requirement that FIs establish and maintain an accurate and up to date AI inventory linked to other relevant inventories such as models, applications, data and third-party services. The inventory becomes a control instrument when it captures attributes that directly support risk decisions. Alongside purpose, scope, model type, data, ownership, lifecycle status and risk ratings, it is useful for material AI to record the level of autonomy and actionability, distinguishing between systems that only recommend actions, those that require human confirmation, and those that execute autonomously, including transaction or value limits where relevant. For AI agents, the inventory can record which internal and external tools an agent can invoke, with what privileges, and under what whitelisted actions rather than relying solely on blacklists. Identifying third party dependencies and update cadence, including provider, model family and version, mode of access such as API, SaaS or on premises, typical update frequency, and whether change notification is contractually required, supports both concentration analysis and change management. Recording whether an AI off mode or manual or alternative system exists for the function, and when such contingency arrangements were last exercised, directly supports operational resilience. **Question 7:** Impact, complexity and reliance are a sound basis for AI risk materiality, and the emphasis on autonomy and availability of alternatives within reliance is well founded. Two further dimensions naturally complement this framework. The first is connectedness or shared

		dependency at portfolio level, which captures the degree to which multiple AI use cases and critical services rely on common models, vendors, data sources or infrastructure. Even a moderate impact AI system can materially increase overall risk if it is one instance of a heavily shared dependency. The second is misuse and adversarial exposure, particularly for Generative AI and agents, where risks include deepfakes, social engineering, prompt injection and data exfiltration. These are already highlighted in MAS' publications on cyber risks associated with Generative AI and deepfakes. Recognising higher connectedness or adversarial exposure as factors that can elevate materiality, even where single use case impact appears moderate, is consistent with the risk profiles MAS has described. **Question 8:** MAS' coverage of AI life cycle controls is comprehensive and aligned with good practice and includes data management, fairness, transparency and explainability, human oversight, third party AI risk, selection, evaluation and testing, technology and cyber security, reproducibility and auditability, monitoring and change management. The risk is that this breadth can turn into a diffuse checklist. For AI used as an integrated part of business processes, a concise Assurance File that maps these expectations to the specific use case and its materiality can ensure clarity without unnecessary duplication. For high materiality AI, the Assurance File core can set out the intended use and explicit boundaries, including what the AI is not to be used for, the materiality assessment including connectedness and adversarial exposure, key data sources and data quality and representativeness constraints, the evaluation approach and acceptance thresholds, including robustness and fairness where relevant, and the key failure modes and mitigants. For Generative AI, those failure modes include hallucinations, toxic or biased content, prompt injection, data leakage and misuse. For agents they include autonomy, uncontrolled tool invocation and goal divergence. The Assurance File can also describe the human oversight design and the accountable owner, third party disclosure gaps and compensating testing and scope restrictions, monitoring metrics, thresholds and escalation responses, and change triggers and rollback plans. For all high materiality AI with customer impact, tested kill switches or AI off modes with documented activation protocols and periodic drills proportionate to criticality, for example at least annually for critical business services, provide strong assurance on the ability to control risk. However, a kill switch that simply breaks the workflow is not sufficient. Firms need to demonstrate degraded mode operations, showing that key services can continue in a controlled, if slower or more manual, state after AI is disabled, and that staff understand the operational playbook for operating under those degraded conditions. For AI agents, minimum Assurance File content can reasonably include explicit transaction or value limits, whitelisted tools and actions, and real time monitoring of tool invocation chains. Given MAS' emphasis on autonomy and goal divergence, high materiality agent use cases also warrant agent specific red teaming, in which independent testers deliberately try to induce undesired goals, tool use and sequences of actions. This is more informative than standard functional testing for uncovering misalignment between stated objectives and the behaviour of agents with tool access. For lower

		materiality AI, it is natural to simplify or omit some of these elements, consistent with the proportionality envisaged in AIRG. MAS already promotes risk-based transparency and explainability. For high complexity models such as large Generative models, mechanistic interpretability is constrained, and emphasis on clarity of inputs, policies, allowed actions, limitations and failure modes, empirical evaluation including robustness, safety, misuse and bias testing, and continuous monitoring of outputs and, for agents, of reasoning paths, tools and actions, is more effective than attempting spurious precision on local explanations. Clear appeal and redress mechanisms for affected customers complete the picture. This is not a relaxation of transparency. Even where detailed mechanistic explanations are not feasible, AIRG can still expect meaningful transparency appropriate to the audience, including clear statements of intended use, key drivers or constraints, limitations and error modes, accountable ownership, and measurable customer outcome monitoring and redress. Where FIs use third party or black box models without full transparency over training data or internals, more extensive testing, constrained deployment scopes, and shorter re validation cycles are a reasonable expectation. For high materiality use cases, where a vendor is unwilling or unable to provide sufficient information to assess conceptual soundness and key design assumptions, may be difficult to conclude it meets AIRG expectations without additional evidence. In such cases, firms should expect to seek alternative vendors or internal builds rather than relying indefinitely on opaque third party models whose risks cannot be meaningfully assessed. **Question 9:** Two additional aspects merit emphasis in AIRG. The first is portfolio concentration and systemic readiness. MAS already points to concentration risk from over reliance on a few Generative AI providers. This can be operationalised by expecting firms to identify their top AI providers by the number and criticality of dependent use cases, map which critical business services depend on each provider, maintain a register of synchronised update channels, and develop and periodically test incident playbooks for scenarios such as major provider outages, regressions or exploits. High level clarity on the expected protocol for coordinated sector wide response when critical vulnerabilities in widely deployed models are identified, including when temporary restriction, heightened monitoring, alternative controls or portfolio level kill switch activation are expected, would further support resilience. The second is pilot discipline and time boxing. Many pilots become de facto production with weak controls. It is helpful if AIRG expects pilots to be explicitly time bounded, for example three to six months, with defined scope in terms of data, users and transactions, and explicit success criteria. Any extension beyond the initial window can then require explicit re approval and a refreshed risk assessment. Controls can be calibrated to scope, with lighter documentation only where limits are tight and monitored. This should be aligned with PDPC expectations on purpose limitation, data minimisation, retention, and transparency for AI enabled decisioning, particularly where personal data is involved. **Question 10:** A twelve month transition period is reasonable, provided expectations are phased and risk based. From the date AIRG is issued, firms can immediately enter an initial three to six

		month period focused on operationalising definitions and Annex based scoping, embedding identification triggers into change and procurement processes, standing up a central AI inventory, and approving a materiality methodology. For Singapore branches and subsidiaries of global FIs, there is clear benefit in leveraging parent entity AI risk frameworks, inventories and materiality assessments, but this cannot substitute for local accountability. A disciplined approach is for the locally accountable person to sign off a documented gap analysis between the global framework and AIRG, covering FEAT, PDPC requirements and other Singapore specific expectations, and to evidence how any gaps are closed in the local implementation. Only where that analysis shows that the group framework meets or exceeds AIRG expectations, and Singapore specific risks are explicitly addressed, should reliance on parent frameworks be regarded as adequate. In the next phase, up to twelve months, firms can reasonably demonstrate that Assurance Files are implemented for high materiality use cases, that monitoring dashboards and escalation are working in practice, that incident playbooks have been written and exercised, and that governance has been integrated into existing risk and committee structures. In the subsequent phase beyond twelve months, firms can extend Assurance File coverage and controls to medium materiality AI, mature portfolio dependency analysis, and roll out capability and training to business and control functions. FIs with complex legacy estates may need longer for full coverage, and that can be accommodated where there are credible interim mitigants and a clear prioritisation towards high materiality use cases, which is naturally supported by the Annex and materiality framework. In parallel, basic AI policies, as envisaged in AIRG, should be implemented commensurate with adoption from day one, even as more advanced elements are phased in.
70	The Hongkong and Shanghai Banking Corporation (Singapore Branch), HSBC Bank (Singapore) Ltd	**Question 1:** HSBC supports the proportionate application of the Guidelines to all FIs. We agree that the framework should be commensurate with each FI's level of AI adoption, maturity and impact, while also considering other factors such as the FI's size, complexity, structure, risk profile and multi-jurisdictional nature. HSBC also supports a proportionate approach to the application of requirements based on the relevance and risk materiality of the AI use cases, systems and models within the FI. A risk-based approach to governance and control application is consistent with HSBC's Group Risk Management Framework. **Question 2:** HSBC acknowledges that the Guidelines are designed to build upon and complement the FEAT principles. However, it appears that the scope of the Guidelines is intentionally broader than the FEAT principles, which are currently limited to AI or data analytics supporting or replacing human decision-making. If so, AI use cases and systems previously outside the scope of the FEAT principles may now be captured under the Guidelines. Greater alignment and clarity on how the FEAT principles and the Guidelines interconnect would be helpful. Additionally, FIs would benefit from further guidance on whether commonly used, non-complex traditional machine learning techniques should be classified as AI or AI models. Applying identical control requirements to these techniques as one would to sophisticated neural networks or generative AI models may not reflect a risk-sensitive approach. We

		welcome MAS' views on whether FIs may exclude such traditional methods from the scope of the Guidelines, so that FIs can better prioritise controls where the actual risks lie. **Question 3:** HSBC agrees that robust oversight by the Board and senior management is essential for effective AI risk management. We request further guidance on whether the reliance permitted in Paragraph 1.3 of the Guidelines—allowing FIs to leverage the AI risk management frameworks of their parent entities—extends to oversight by global senior committees or Boards (that have formal responsibility for AI governance) of such AI risk management frameworks. If so, this would make the allowance in Paragraph 1.3 much more meaningful and enable global FIs operating in Singapore to achieve greater scale and efficiency in AI governance. If MAS is supportive, we welcome guidance on what types of evidence would be considered sufficient to demonstrate effective oversight at a local level. We would assert that such oversight could be substantiated through existing governance structures established under the Group Risk Management Framework, rather than requiring a local specialist committee or forum. We also welcome clarification in paragraphs 2.2 to 2.4 of the Guidelines to more clearly distinguish between the roles and responsibilities of the Board and those of senior management, as the current language refers to "the Board and senior management" collectively without distinction. **Question 4:** HSBC does not treat AI as a discreet risk type. AI both affects and is affected by various types of risk throughout its development lifecycle – for example, model risk, data risk, or cyber risks. Effective management of these risks is established through adherence to our Group Risk Management Framework and associated risk policies and procedures. Therefore, we are supportive of the less centralised approach, where risks impacted by AI are assigned to and managed through existing risk management functions. HSBC employs a globally standardized approach to AI governance, which enables holistic coverage of risks and consistency, while allowing autonomy in execution across our global businesses and functions. As part of this approach, HSBC has established dedicated cross-functional committees and councils to oversee AI at both the Group and business level. Certain business level councils have joint responsibility for overseeing AI within their line of business as well as across the markets in their region, closely aligning with the HSBC Group governance structure. HSBC would welcome guidance from MAS on the supervisory expectations in leveraging dedicated cross-function AI committees that operate at the global business level, with regional and local oversight. **Question 5:** HSBC is supportive of the proposal and have no comments.

		Question 6: HSBC supports the establishment and maintenance of an accurate and up-to-date inventory of AI use cases, systems or models. However, as noted in our response to Question 2, we would appreciate guidance on whether this inventory is expected to include non-complex, traditional machine learning techniques. The attributes tracked for such techniques would be significantly fewer than those required for generative AI or agentic AI use cases. Additionally, it would be helpful to understand what aspects of the AI risk materiality assessments MAS expects to be included in the inventory. Further comments on the AI risk materiality assessment are provided in our response to Question 7. **Question 7:** HSBC supports the inclusion of the three proposed dimensions—impact, complexity, and reliance—as a sound basis for assessing the different risks associated with the use of AI. Our views are similar to those reflected in the Mindforge AI Risk Management: Executive Handbook (November 2025) (Mindforge), where financial institutions should rely on their existing risk frameworks and enhance them, as appropriate, with any new considerations stemming from the use of AI. HSBC does not treat AI as a discrete risk type. The risks associated with AI are multi-faceted and manifest across various risk types already present in our risk taxonomy, which is underpinned by our suite of policies, procedures and controls. For example, our model risk tiering and IT system tiering frameworks provide structured approaches to assess and manage these risks. Therefore, at HSBC, we review any use of AI against our existing Global Risk Framework and domain-specific risk methodologies – including data, technology/resilience/BCP, model, cyber, and third-party risk. This approach provides the structure, taxonomies and tools to mitigate and manage these risks effectively. Where heightened risk is identified, our methodologies require enhanced controls, ongoing monitoring, and, where appropriate, increased scrutiny and escalation protocols tailored to the specific risk domain. This ensures that elevated risks are managed in a manner consistent with their nature and impact. To ensure our risk taxonomy sufficiently addresses risks associated with the use of AI, we have a process to identify AI risk events and map them to the relevant domain-specific risks and supporting controls in our taxonomy. We then uplift these methodologies and controls to ensure they remain robust and fit for purpose as AI technologies evolve. HSBC is therefore supportive of keeping the AI risk materiality assessment framework flexible, allowing alignment with existing risk assessment tools. This approach ensures that risk assessments remain proportionate and relevant to the actual risks presented by each use case, in line with the industry practices reflected in Mindforge. **Question 8:** HSBC supports the requirement to implement robust controls across the entire AI lifecycle. We agree with the key principles outlined in Section 6, Paragraphs 4.5 to 4.26 of the Guidelines.

		Further guidance regarding MAS' expectations for the selection of AI algorithms and features for specific AI use cases would be beneficial. We support the requirement for FIs to select algorithms and models that are fit for purpose. However, the expectation that such selections should be supported by theory, research, or accepted industry practice—and potentially incorporate reviews by domain experts or users for each use case—is a high threshold if applied universally. This may not always be practicable for every AI use case, especially given the diversity and pace of AI use cases and their implementation in FIs. Additionally, regarding contingency plans and fallback options for high-risk AI use cases, we note that such plans may not need to be AI-specific. As AI models and systems are part of an FI's broader IT infrastructure, these considerations are already an integral part of an FI's IT framework and business continuity planning (BCP). We believe that contingency plans and fallback options may best be managed within existing IT and BCP frameworks, rather than requiring separate AI-specific measures in all cases. We also recognise the critical importance of post-deployment monitoring and review as part of the AI life cycle. Further guidance would be appreciated on whether such monitoring is expected to be conducted at the individual AI use case level, or if it must be performed as an aggregate, risk portfolio analysis. Given the diverse nature of AI use cases, we believe that monitoring at the individual use case level would be more effective, enabling targeted responses to specific issues and ensuring the unique characteristics of each use case are appropriately addressed. **Question 9:** HSBC has no additional comments. **Question 10:** HSBC is supportive of the proposed 12-month transition period, subject to the clarifications provided in response to this consultation paper. We would welcome guidance on when the transition period will commence. It would be preferable for the transition period to begin after the supporting AI Risk Management Handbook have been reviewed and published, to ensure there is sufficient clarity regarding implementation requirements.
71	Thunes	**Question 1:** MAS can further clarify or provide examples if aggregation of multiple low-risk "assistive" tools could collectively trigger or become material at organizational level, which provides further guidance on "basic policy" and "integrated processes" route. **Question 2:** MAS can provide examples of level or definition of "autonomy" threshold for AI Agents for organizations to assess their implementation. **Question 3:** We support this on expectation for the Board to have a broad understanding of AI is vital and similar to ISO 42001. **Question 4:**

		We support this approach that a dedicated committee is necessary and the assessment should be on overall potential risk and impact which is similar to ISO 42001 **Question 5:** We support this as consistent identification is important to ensure robust systems and prevent "shadow AI" across the organization. This is critical to reduce risk exposure for organization. **Question 6:** MAS can clarify if use of tools that serve as automated model registry is permitted to meet or support this requirement. Automated model registry tools provide ease of use and upkeep while allowing auditability. **Question 7:** **Question 8:** We support the idea of "kill switches" as safety mechanisms for high-risk AI use cases to prevent cascading financial loss and support the idea of implementing human oversight to prevent automation bias by having a human-in-the-loop approach. **Question 9:** **Question 10:**
72	Transamerica Life Bermuda	**Question 1:** No further concern **Question 2:** No further concern **Question 3:** No further concern **Question 4:** No further concern **Question 5:** No further concern **Question 6:** No further concern **Question 7:** While the consultation paper contains expectations on third party AI risk management, it does not clearly define on how an FI should technically review AI model or service developed by third parties — especially when AI is cloud hosted, closed source, or lacks transparency. (e.g. What if source code or model internals are unavailable)? Also, we would like to seek for more guidance and examples on how to evaluate large cloud AI

		providers with limited disclosure, like Google, AWS, Azure? Should this be classified as "transparency not available"? **Question 8:** No further concern **Question 9:** No further concern **Question 10:** No further concern
73	Transparentl y Pte Ltd	**Question 1:** Transparently.AI fully supports the Monetary Authority of Singapore's (MAS) efforts to establish a robust and proportionate AI risk management framework through the proposed Guidelines on Artificial Intelligence Risk Management (AIRG). We recognise the need for clear standards to ensure safety, ethics, and consumer protection. However, as a leading AI provider serving the financial sector, we observe significant hesitancy among Financial Institutions (FIs) regarding AI adoption, driven by uncertainty and fear of regulatory missteps. Our key concern is that the current phrasing and structure of the Guidelines, particularly regarding mandatory governance structures and highly prescriptive controls for all identified AI usage, may create a "regulatory chill effect." This could inadvertently halt innovation and widen the technology gap, especially for FIs that are currently nascent in their AI adoption journey. We propose an "Enablement-Focused" approach that strongly differentiates between low-risk/assistive AI and integrated/high-risk AI, offering simplified compliance pathways for the former to encourage experimentation and confidence building. The primary objective of the Guidelines must be to foster Responsible Adoption, not merely prevent risk. Currently, the administrative burden implied by the requirements (Oversight, Identification, Inventory, Assessments) could lead risk-averse FIs to simply postpone or abandon all but mandatory AI projects. Recommendation A: Implement a Clear De-Minimis/Safe Harbour Provision MAS has correctly established proportionality based on "integrated part of business processes" (Annex, para 2 & 4). We recommend strengthening this concept by providing clearer de-minimis thresholds for non-integrated, low-risk use cases (e.g., those assisting only internal staff with non-critical tasks). Proposal: Use cases meeting the following criteria are only subject to the Basic Policies (Annex, para 5) and are explicitly exempted from the full Section 2 (Oversight) and Section 3 (Key Systems) requirements: Low Impact: No direct customer or third-party interaction (internal function only). Human-in-the-Loop: Requires mandatory human review, validation, and override capability on 100% of outputs used for critical decisions.

		No Material Financial Impact: The immediate failure of the system would result in a financial loss exposure of less than 0.01% of the FI's annual operating expenditure (OPEX). This creates a low-friction zone for FIs to build comfort and capability. Recommendation B: Focus on Outcomes-Based, Rather Than Process-Based, Controls The AIRG is comprehensive in listing Life Cycle Controls (Section 4). To support diverse and fast-moving AI technologies, the MAS should emphasize the risk outcomes (e.g., bias mitigation, system stability) rather than mandating specific control processes (e.g., a specific type of stress test). This allows FIs and technology providers to adopt best practices that match the specific technology they employ (e.g., different controls for traditional ML vs. Generative AI). While proportionality is essential, the distinction between "basic policies" and the full framework must be clearer. We recommend MAS provide quantifiable examples of when an FI transitions from the "Basic Policies" tier (Annex, para 5) to the full framework requirements (Sections 2-5). This removes ambiguity and reduces the fear of under-complying. **Question 2:** The definition of AI (para 1.2) is appropriately broad. However, to avoid the "chill effect," FIs should be assured that the full compliance and documentation burden (Inventory, Cross-Functional Committee, formal Validation) is strictly and solely applied to use cases classified as Material based on the Risk Materiality Assessment (Section 3). **Question 3:** Mandatory establishment of a Dedicated CFCC (para 2.4) when overall risk is "deemed material" creates an administrative hurdle that can deter adoption. For FIs with low AI maturity, the initial oversight should be handled by an existing management committee (e.g., a Technology Risk Committee). The CFCC should only become mandatory when the number of high-risk use cases reaches a predefined volume or scale (e.g., >3 material use cases, or a specific financial exposure threshold), ensuring the governance structure scales with adoption. **Question 4:** Mandatory establishment of a Dedicated CFCC (para 2.4) when overall risk is "deemed material" creates an administrative hurdle that can deter adoption. For FIs with low AI maturity, the initial oversight should be handled by an existing management committee (e.g., a Technology Risk Committee). The CFCC should only become mandatory when the number of high-risk use cases reaches a predefined volume or scale (e.g., >3 material use cases, or a specific financial exposure threshold), ensuring the governance structure scales with adoption. **Question 5:** We agree that identification (Q5) and inventory (Q6) are critical. The proposed systems (para 3.2-3.7) should focus resources where risk is highest. The rigor required for documentation and maintenance of the Inventory (para 3.5) should be directly proportionate to the assessed risk materiality. A low-risk, non-integrated internal tool

		should require minimal inventory data, while a high-risk underwriting model requires maximal detail. **Question 6:** We agree that identification (Q5) and inventory (Q6) are critical. The proposed systems (para 3.2-3.7) should focus resources where risk is highest. The rigor required for documentation and maintenance of the Inventory (para 3.5) should be directly proportionate to the assessed risk materiality. A low-risk, non-integrated internal tool should require minimal inventory data, while a high-risk underwriting model requires maximal detail. **Question 7:** The three proposed dimensions are comprehensive and appropriate for the financial sector. To aid FIs in consistent self-assessment and reduce regulatory anxiety, MAS should provide a suggested tiered scoring model (e.g., a simple 1-5 score for each dimension) to clearly guide the classification of a use case as High, Medium, or Low Risk Materiality. **Question 8:** We stress the need for flexibility. For example, regarding Transparency and Explainability (para 4.6), the control should be the outcome—e.g., "The FI must be able to provide a clear, tailored explanation of the decision to the affected customer." The method used (e.g., LIME, SHAP, model simplification) should be left to the FI/developer, allowing for innovation and use of complex, high-performing models. **Question 9:** To ensure controls facilitate, rather than impede, adoption, we offer the following specific refinements to Section 4: Transparency and Explainability (Para 4.6, 4.7): Focus on Use-Case Specific Explanations: Clarify that explainability should be commensurate with the audience and context. For a customer, a non-technical explanation of the key drivers is sufficient. The requirement for internal technical explainability (e.g., full model feature importance) should be reserved for high-risk models only, preventing the over-simplification of high-performing AI simply to meet a generic explainability standard. Human Oversight (Para 4.10): Mitigating Automation Bias: The Guidelines correctly identify automation bias. We recommend explicitly requiring FIs to implement structured feedback loops where human reviewers are required to justify overrides of AI-generated decisions. This mandates active engagement and reduces passive reliance, ensuring the human remains "in control" or "over the loop" without increasing the administrative overhead of continuous monitoring. Third-Party AI Management (Para 4.11):

		Standardising Disclosure Expectations: Third-party providers (like our company) want to comply. We urge MAS to work with the industry to define standardised minimum disclosure requirements for third-party AI, particularly concerning data provenance, high-level architecture (for complexity assessment), and known limitations/failure modes (hallucinations, bias testing results). This clarity reduces the testing burden on individual FIs (compensatory testing, para 4.11) and promotes market confidence. Conclusion: We believe the AIRG framework is a strong foundation for responsible AI. By adopting an Enablement-Focused mindset—specifically by providing a clear "Assistive AI Safe Harbour," promoting flexible, outcomes-based controls, and allowing governance structures to scale with maturity—MAS can successfully guide the industry towards both safety and accelerated innovation. We welcome the opportunity to further discuss these recommendations. **Question 10:** 12 months (para 4.7) may be adequate for FIs with high AI maturity. However, FIs currently operating under the "Basic Policies" tier will need significantly longer (e.g., 18-24 months) to establish the full Section 2 and 3 frameworks (Oversight, Inventory, Risk Assessments). A tiered approach ensures a practical pathway for everyone.
74	UAPF/Algomation	**Question 1:** I support proportionate application across all FIs. The Unified Algorithmic Process Format (UAPF) provides concrete implementation mechanisms using international standards (BPMN 2.0, DMN 1.3, CMMN 1.1). Full proposal: https://www.notion.so/Proposal-for-Algorithmic-Process-Governance-in-Financial-AI-Systems-2f688e97feb2801ca82adfec500c9a95 Key implementation approach: - Small FIs: Simple process documentation as starting point - Large FIs: Advanced automation with full governance stack - Scales naturally with organizational complexity - Leverages existing international standards (OMG-maintained) - Open-source specification enabling vendor-neutral adoption **Question 2:** **Question 3:** Board and senior management require structured tools for oversight. UAPF enables this through: Full proposal: https://www.notion.so/Proposal-for-Algorithmic-Process-Governance-in-Financial-AI-Systems-2f688e97feb2801ca82adfec500c9a95 Implementation mechanisms: - Machine-readable process specifications (BPMN) for board-level visibility

		- Audit trails from objectives → AI outputs → outcomes - Version-controlled change management - Clear accountability mapping **Question 4:** **Question 5:** FIs should adopt structured process modeling standards combined with UAPF for consistent AI identification. Full proposal: https://www.notion.so/Proposal-for-Algorithmic-Process-Governance-in-Financial-AI-Systems-2f688e97feb2801ca82adfec500c9a95 Technical approach: - BPMN 2.0 for process flows - DMN 1.3 for decision logic - CMMN 1.1 for case management - Automated tracking across organization - Machine-readable specifications **Question 6:** Agreed. AI inventories should be maintained at the process level, not just model level. Process registries (similar to model registries but for business processes) should store version-controlled process definitions, link processes to AI model inventories, track ownership, and enable automated compliance checking. The UAPF approach enables this through standardized, machine-readable process specifications. **Question 7:** The proposed dimensions (impact, complexity, reliance) are appropriate. However, risk materiality should be evaluated within process context, not just for isolated models. This means: (1) Impact: assess consequences within end-to-end processes including cascading effects; (2) Complexity: evaluate process integration complexity and interdependencies; (3) Reliance: measure organizational dependency on AI-enabled processes. UAPF enables structured assessment by providing clear process boundaries, documented decision logic, and traceable impact pathways. **Question 8:** I support the proposed lifecycle controls and recommend implementing them through process-level governance: (1) Data Management: document flows using BPMN data objects; (2) Transparency/Explainability: process models provide business logic explainability complementing model-level explanations; (3) Human Oversight: define checkpoints within process flows using BPMN user tasks; (4) Third-Party Risk: map external AI services within process architecture; (5) Monitoring/Change Management: version-controlled process specifications enable systematic tracking, impact analysis, and audit trails. **Question 9:**

		I recommend MAS consider Algorithmic Process Governance as a foundational framework for implementing these Guidelines. COMPLETE PROPOSAL: https://www.notion.so/Proposal-for-Algorithmic-Process-Governance-in-Financial-AI-Systems-2f688e97feb2801ca82adfec500c9a95 This submission proposes: 1. PROCESS-LEVEL GOVERNANCE - AI models operate within business processes - Risk management must address process context, not just isolated models - Process registries complement model registries 2. INTERNATIONAL STANDARDS - BPMN 2.0 for process flows - DMN 1.3 for decision logic - CMMN 1.1 for case management - OMG-maintained, widely adopted 3. UAPF SPECIFICATION - Open-source framework: https://uapf.dev - GitHub: https://github.com/UAPFormat - Machine-readable process documentation - Enables automated compliance checking 4. IMPLEMENTATION BENEFITS - Scalable across FI sizes - Vendor-neutral approach - Cost-effective implementation - RegTech enablement 5. ALIGNMENT - EU AI Act technical documentation requirements - NIST AI RMF governance principles - Basel Committee operational resilience - FSB AI/ML frameworks **Question 10:** 12 months is appropriate, allowing FIs to: Phase 1: Document highest-risk AI processes; Phase 2: Establish process registries and governance workflows; Phase 3: Expand to medium-risk processes; Phase 4: Achieve comprehensive coverage. This phased approach aligns with proportionate implementation principles.
75	United Overseas Bank Ltd	**Question 1:** We appreciate and support the principle of proportionality in the implementation of these Guidelines. However, we encourage MAS to adopt a pragmatic approach in practice. Should there be differing interpretations of what constitutes a proportionate approach for a bank, the bank should be given the benefit of the doubt and not be immediately faulted for non-compliance.

		Instead, MAS could allow the bank reasonable time to address any identified gaps and align with supervisory expectations. This approach will help ensure that implementation remains practical and constructive, while still upholding the intent of the Guidelines. **Question 2:** We believe it is never the intent of MAS to limit the scope of these Guidelines to only AI models, systems, and use cases. The comprehensive approach outlined in the consultation paper clearly aims to address AI risk management across the entire bank. To avoid any potential misinterpretation, we suggest that MAS consider rewording the scope to explicitly cover end-to-end expectations for AI risk management. This should include not only models, systems, and use cases, but also supporting infrastructure, governance, processes, people, and all other relevant elements necessary for effective AI risk management throughout the lifecycle. Clearer wording of the scope will help ensure consistent understanding and implementation across the industry, fully reflecting MAS's intent for a holistic approach. In addition, we seek clarification on whether models developed using a fully judgmental approach – where output scores are determined solely by inputs from subject matter experts, without any learning or inference from input data – are excluded from the definition of AI under these Guidelines. Our understanding is that such models are analogous to "tools whose outputs are solely based on predefined programming logic or rules," as referenced in Paragraph 1.2(d), and therefore would not fall within the scope of AI as defined in the consultation paper. **Question 3:** We agree with MAS's proposal regarding the responsibilities of the Board and senior management in overseeing AI risk management. We would like to emphasize that the expectation for Board and senior management oversight of AI risk management should be no different from their existing responsibilities for overall risk management. AI risk management should be embedded within the bank's overall enterprise risk management approach. In this regard, banks should leverage their established risk management frameworks to address AI-related risks, ensuring consistency and integration across all risk domains. **Question 4:** We agree with MAS's recommendation that a dedicated cross-functional committee should be established where the overall AI risk exposure is high. Such a committee can provide effective oversight and coordination across functions. To assess overall AI risk exposure at the organizational level, we recommend a structured approach. The approach should align with the extent to which AI is deeply embedded across all areas of the organization and becomes integral to its operations and value delivery. AI risk exposure is deemed to be high where AI is not just used as an occasional tool, but as a core component of business strategy and daily processes.

		Question 5: We agree with MAS's proposal regarding the establishment of clear definitions, criteria, and processes—supported by robust systems—to facilitate the consistent identification of AI usage across all relevant business and functional areas. We would also like to seek MAS' clarification on what constitutes "robust systems" for identifying AI. Is MAS referring to a system to inventorise and capture all AI usage across the Bank? It would be helpful if MAS can elaborate with some examples. **Question 6:** While we agree that maintaining an AI inventory is important, we would like to highlight two points for MAS's consideration: 1. Flexibility in Inventory Management: The manner in which the AI inventory is maintained and managed should be left to the discretion of each bank. It need not necessarily be maintained in a dedicated system; an Excel spreadsheet or other suitable format should suffice, as long as it meets the intended objectives. 2. Scope and Minimum Requirements: It would be helpful if MAS could set out its minimal expectations regarding the information that must be maintained in the inventory. The entire model documentation should not be required; rather, the focus should be on key attributes necessary for effective oversight and risk management. Additionally, we suggest that the requirement to maintain an AI inventory should apply only to AI use cases, systems, or models assessed as medium or high risk. Low risk use cases should not be mandatorily included in the inventory. **Question 7:** 1) Two-step approach for risk materiality assessment We recommend adopting a two-step approach for assessing the risk materiality of AI use cases, systems, or models: 1. Step 1 – Impact Assessment: The first step should be to determine the potential impact if the AI model, use case, or system were to fail or perform poorly, as well as to consider the sensitivity of the data being used. If it is established that the impact is low, then further assessment of complexity and reliance becomes less critical. 2. Step 2 – Complexity and Reliance (Conditional): Only if the impact is assessed as medium or high should the assessment proceed to consider the complexity and reliance dimensions. In such cases, higher complexity or greater reliance would indicate higher risk and warrant tighter governance and controls. In summary, if the impact is low, the complexity and reliance factors should not materially affect the risk assessment. This approach helps ensure that resources are focused on higher-impact AI applications, while maintaining proportionality in risk management. 2) Clarification on inherent risks and residual risks

		Paragraph 3.8 of the Guidelines states that risk materiality assessments should consider both inherent and residual risk. We note that Section 4 references AI use case, system or model assessed as "high risk" in several paragraphs (4.4, 4.19, 4.20, 4.23, and 4.24). We seek clarification on whether these references pertain to inherent risk or residual risk. Additionally, we recommend that the Guidelines provide more explicit guidance, or examples, on how inherent and residual risk should be assessed, specifically with respect to the three risk dimensions outlined – impact, complexity, and reliance. **Question 8:** We agree with MAS's proposals regarding the standards, processes, and controls that should be applied across the entire AI lifecycle. However, we would like to seek further clarification on the requirements set out in Paragraph 4.22 of the Proposed Guidelines which state that FI should "conduct vulnerability assessment, penetration testing, and red teaming across AI systems". • To better understand the supervisory expectations, we would like to clarify on the intended scope of "vulnerability assessment, penetration testing and red teaming" in the context of AI. (a) Whether the proposed guidelines distinguish between AI Red Teaming (e.g., attempting to jailbreak a model via prompt engineering to bypass safety filters), and Red Teaming (e.g., attacking API endpoints and plugins connected to the model). We would appreciate MAS' clarification on whether both AI specific and infrastructure-based red teaming are required to be performed for this requirement and whether the approach should be manual, automated or combination of both. (b) On the requirement to perform vulnerability assessments and penetration testing, we would like to seek clarification on whether the guidelines' scope is referring strictly to traditional infrastructure scanning (e.g., container security, CVE scanning), or AI-specific artifact scanning (e.g., scanning model weights/files for embedded malware or pickle serialization attacks). Given the evolving nature of AI specific vulnerabilities, we would seek MAS guidance on AI security testing and AI red teaming. Such standards would support FIs in building robust and consistent practices to identify vulnerabilities specific to AI models and systems. Lastly, we would like to propose for consideration, referencing "adversarial threats" mentioned in Paragraph 4.22 (b) to established taxonomies such as MITRE ATLAS (Adversarial Threat Landscape for AI Systems). This will eliminate any ambiguity or misalignment in the interpretation across the industry. **Question 9:** No comment. **Question 10:** Given that these are guidelines, we believe MAS should allow banks the flexibility to implement the recommendations on a proportionate basis. In other words, the proposed

		one-year transition period should not be a hard line in the sand covering all aspects of the recommendations. Banks should be able to prioritize and phase in implementation according to the materiality and complexity of their AI use cases, ensuring a practical and effective adoption of the Guidelines.
76	UOB Kay Hian Private Limited	**Question 1:** **Question 2:** **Question 3:** **Question 4:** proposing that such a committee need not be standalone body, and may be integrated into existing governance committee **Question 5:** **Question 6:** **Question 7:** **Question 8:** **Question 9:** provide guidance on the interaction between the AI Guidelines and existing MAS requirements (eg TRM, Outsourcing) **Question 10:**
77	US-ASEAN Business Council	**Question 1:** FIs in Singapore that are branches or subsidiaries with parent entities in other jurisdictions may leverage the AI risk management frameworks of their parent entities, as long as such frameworks meet the expectations set out in these Guidelines. While this approach aims for consistency, its breadth of application may create disproportionate compliance burdens for entities with minimal consumer-facing AI use. Wholesale or institutional-focused entities typically deploy AI for internal analytics or operational efficiency, posing negligible direct consumer harm. Recommendation: Limit application of the Guidelines to allow risk-based proportionality for other FIs, focusing on risk materiality, and regulatory impact rather than blanket inclusion. This targeted approach ensures MAS achieves its consumer protection objectives while avoiding unnecessary compliance costs for low-risk entities. ==

		Do the MAS guidelines apply only to AI use cases supporting processes for clients booked in Singapore, or do they also cover scenarios where a client is booked in Hong Kong but managed or advised from Singapore? == Most third-party services are expected to incorporate AI tooling into non-specific services, which could result in a massive expansion of the scope of applicability to third-party services. From a practical and operational perspective, such third-party services may not be able to fully comply with the type of governance and oversight set out in the draft guidelines. Recommendation: The MAS guidelines would benefit from further clarity regarding their application to third-party AI providers and services. The MAS guidelines should align with the MAS Technology Risk Management Guidelines and other relevant Notices (e.g., Cyber Hygiene Notice) and Guidelines (e.g., Guidelines on Outsourcing) to ensure integration and cohesion with existing frameworks, rather than repeating high-level concepts. This would clarify how AI-specific controls build on established frameworks. We seek confirmation that the proposed guidelines would be applied in a harmonized manner, particularly for activity and transactions that are international in nature. For example, for a client managed in Hong Kong and advised from Singapore, the bank would be subject to Hong Kong regulations and would seek clarity on whether processes could be leveraged to comply with Hong Kong regulatory requirements for the purposes of fulfilling MAS's expectations. == The industry supports MAS's proportional approach to AI risk management. The extent of oversight, as well as the design and implementation of AI risk management systems, policies, and procedures, should be proportionate to the size and nature of an FI's activities, its use of AI, and its overall risk profile. This flexibility is essential to ensure requirements are both effective and practical, allowing FIs to calibrate controls to their unique business models and risk exposures. This focus on proportionality reflects a global consensus: leading AI frameworks in the EU, UK, US, Canada, and Japan all prioritize governing risks rather than the technology itself. By extending oversight to the materiality of AI use, organizations can maximize the benefits of AI while ensuring meaningful safeguards. Importantly, this risk-based approach enables businesses to fulfill multiple regulatory regimes simultaneously, reducing duplication and supporting innovation across markets. Where firms deploy third-party AI technologies without material modification to the underlying model or its intended purpose, reliance on vendor documentation, attestations, and recognized international standards may be appropriate for oversight and inventory purposes. Contextual configurations or usage-level adaptations (e.g., prompt design, parameter tuning within vendor-provided guardrails, workflow integration, or applying the model to enterprise use cases) should not, on their own, trigger full

		re-validation or lifecycle controls, provided they do not introduce new or higher-risk outcomes. Recommendation: Clear guidance on how proportionality applies—especially for FIs leveraging thirdparty technology providers—will help reduce compliance duplication and operational burden while maintaining robust risk management. For example, MAS could allow FIs to rely on formal attestations from technology providers, including recognition of international standards, as evidence of compliance. MAS is encouraged to explicitly recognize international certifications and standards (e.g., ISO/IEC 42001, NIST AI Risk Management Framework, EU CE marking for AI) as valid evidence of compliance with proportionality requirements. This will facilitate cross-border compliance and reduce duplication for global organizations. Clarifying this distinction will help ensure that obligations remain proportionate, outcomes focused, and do not unintentionally discourage responsible customization for legitimate business use cases. == Although AI may be "integrated" into business processes, this does not necessarily imply that its application constitutes a higher-risk use of AI. There can be scenarios where an AI model, system, or use case is embedded within business operations yet presents minimal risk—such as basic automation, nongenerative functionalities, extractive tasks, or structured data annotation. The current approach to defining "integration as part of the business" does not appear to provide an exemption for these low-risk applications, which seems inconsistent with the broader principle of a proportional application of the MAS Guidelines. We also support MAS's recognition that low-risk AI applications should not be subject to the same stringent requirements as high-risk applications that directly affect customer outcomes. AI tools serving assistive functions—such as copilots for writing, coding assistants, internal document summarization tools, or embedded systems for operational efficiency—present fundamentally different risk profiles compared to AI systems that make autonomous decisions impacting customers or business results. These assistive tools augment human capabilities rather than replace human judgment, and their outputs remain subject to human review and approval before any consequential action is taken. Adopting a proportionate approach enables financial institutions to responsibly experiment with and implement beneficial AI technologies without being deterred by disproportionate compliance costs. The use of AI is not new to the financial sector, and some FIs have been exploring and adopting the use of AI and machine learning (ML) applications for more than a decade. For example, some retail payment functions rely on statistical and predictive models that use historical data to generate probabilistic outputs such as scores, likelihoods, or forecasts. These models do not autonomously generate new content or reasoning, but instead optimize towards objectives defined by humans. In retail payments, such models are widely used for fraud risk scoring in card-notpresent transactions, creditworthiness assessments for instalment-based payment products, and forecasting liquidity or

		settlement needs for merchants. These systems are generally well understood, auditable, and governable using established model risk management practices. It is also important to recognize that while generative AI introduces novel risks, such as hallucination or inappropriate data disclosure, it is often separated from core transaction execution and decisioning functions, and therefore presents a different risk profile from automated financial decision-making systems. In payments contexts, generative AI is commonly used for customer service chatbots explaining transaction outcomes, automated drafting of dispute communications, or summarizing internal compliance documentation. Rather than applying the proposed AI risk management guidelines to broadly defined AI models, systems, or use cases, MAS may wish to limit the applicability of the proposed Guidelines to higher risk generative AI use cases so as to not hamper the innovative development of AI/ML tools to deliver better outcomes for customers. An approach that considers the degree of autonomy, customer impact, and decision-making authority supports proportionate guidelines that safeguard trust and stability while enabling innovation. Recommendation: We encourage MAS to provide practical guidance and illustrative examples to help FIs and technology providers distinguish between core AI systems, high-impact use cases, and more peripheral or assistive AI functionalities. This will help ensure that identification efforts are focused and proportionate, and that compliance resources are directed to areas of greatest risk. We recommend that the Guidelines should not apply to AI tools which serve assistive functions or do not determine automated customer outcomes. The risks posed by such assistive tools are in any event governed by other FI-related regulations such as the MAS Technology Risk Management guidelines. Further to the above, we recommend that AI use cases regarded as an "integrated part of business processes", as provided in the Annex, be refined to include consideration of risk proportionality, particularly to distinguish between high risk (e.g., determining customer outcomes) and low risk (e.g., internal productivity/assistive applications). Some of the examples noted under paragraph 4 of the Annex include AI use cases which are internal productivity tools, such as AI legal contract review tools and internal IT helpdesk chatbots. Such clarification would provide greater certainty to FIs on the scope of AI tools that can be excluded from the proposed Guidelines and ensure that supervisory focus remains on high-risk AI tools. We also note that while core principles such as fairness, accountability, and transparency are enduring, the specific controls, processes, and assurance mechanisms used to operationalize these principles should remain adaptive. Allowing implementation approaches to evolve alongside technological developments will help ensure that AI governance remains both effective and futureproof, without constraining innovation through static or overly prescriptive requirements. **Question 2:** The Guidelines adopt a broad definition of artificial intelligence and treat "AI use cases," "systems," and "models" as interchangeable for purposes of risk assessment and lifecycle controls. This conflation introduces conceptual ambiguity regarding the practical

		application of the framework. Moreover, it may obscure the fundamental principle that risk is driven by the manner in which AI is deployed, rather than by the inherent characteristics of the model or system. This approach risks diverting the framework from a genuinely risk-based and proportionate methodology and may create misalignment with regulatory practices in other jurisdictions. Recommendation: We recommend that the Guidelines avoid the interchangeable use of “use case”, “model”, and “system” in the AI definition. == We support MAS's broad and forward-looking scope for the Guidelines, which covers AI use cases, systems, and models—including emerging technologies such as Generative AI and agentic AI. This comprehensive approach is appropriate given the rapid evolution of AI and its diverse applications in the financial sector. MAS's recognition that FIs may leverage parent or group-level AI risk management frameworks, provided these meet local expectations, is also commendable. This flexibility is important for global organizations and their technology partners, as it helps streamline compliance and reduce the burden on institutions that rely on shared resources within a global operating model. Recommendation: MAS is encouraged to continue providing practical examples and clarifications on the boundaries of what constitutes an AI use case, system, or model—particularly in distinguishing between core, high-impact AI and more assistive or peripheral AI functionalities. This will help FIs, and technology providers apply the Guidelines consistently and avoid unnecessary compliance burden for low-risk or purely assistive AI tools. In addition, MAS should continue aligning the scope of its Guidelines with international standards and definitions, to facilitate cross-border compliance and operational efficiency for both FIs and third-party providers. **Question 3:** We agree with the importance of establishing robust frameworks, policies, and procedures to support AI risk management in FIs. Clear governance and accountability are essential for effective risk management, especially as AI becomes more integrated into financial services. Effective governance can be achieved through the expertise and accountability mechanisms of the Boards of Directors and Senior Management whose roles and responsibilities have been clearly set out in existing MAS Guidelines on Risk Management Practices, Technology Risk Management, Outsourcing and Individual Accountability and Conduct. The extent of oversight should be proportionate to the materiality of AI risk exposure, and that dedicated cross-functional committees may be appropriate where AI risk is significant.

		Requiring a dedicated cross-functional committee to govern AI risk exposure may lead to duplicative internal mechanisms and overly prescriptive policies/processes. We encourage a focus on governance outcomes, such as clarity of accountability, effective escalation mechanisms, and appropriate senior management oversight. This approach would allow FIs to design governance arrangements that are effective and address AI-related risks within their existing operating models while meeting MAS's supervisory objectives. This flexibility also allows FIs to tailor governance structures to their business models and risk profiles, while maintaining high standards of risk management and accountability. Recommendation: We propose that MAS provide flexibility in governance structures for AI oversight, without the need for a dedicated cross-functional committee to be set up to govern AI risk exposure. In addition, we recommend MAS adopt a risk-based approach to governance. Regulation should focus on higher-risk contexts and use cases, allowing organizations to continue promoting innovation while instilling public trust. By concentrating oversight on high-risk applications, FIs can allocate resources effectively and avoid unnecessary burden for low-risk or assistive AI tools. In general, Model Risk management frameworks should apply this principle by assessing risk across multiple dimensions—including context, use case, and potential impact—and tailoring governance and review requirements based on the assigned risk rating. We encourage MAS to provide practical guidance or illustrative criteria for assessing when overall AI risk exposure is "material." This could include factors such as the scale and criticality of AI use cases, potential impact on customers and operations, reliance on third-party AI providers, and the complexity or novelty of deployed AI systems. Such clarity will help FIs calibrate their governance structures appropriately and ensure that oversight is commensurate with risk. To support consistent application, MAS could provide illustrative thresholds or rubrics for high, medium, and low risk AI use cases, enabling proportionate governance, examples, and edge cases. Moreover, MAS may consider leveraging existing resources such as IMDA's Compendium of Use Cases for AI Governance. A similar reference, with a focus on financial services, could complement MAS guidance by providing practical examples, clear exclusion, and risk-tier rubrics, helping FIs apply governance consistently and proportionately. This risk-based approach, grounded in real-world guidance, will ensure that regulatory expectations are aligned with materiality, promote innovation responsibly, and strengthen trust in AI adoption across the financial sector. == Paragraph 2.3a:

		The current phrasing, “identify and assess all relevant AI risks,” appears overly broad and may impose unnecessary obligations. We respectfully suggest refining this to “identify and assess all material AI risks, == Paragraph 2.5b: In general, it appears that the guidelines conflate responsibilities for the board and senior management. We think that the Board should be responsible for setting strategy and risk appetite, senior management responsible for implementing strategy and overseeing day to day operations, ensuring they are executed within the approved risk appetite. Recommendation: We recommend MAS not require explicit Board-level responsibility for AI. The Board already has ultimate accountability for enterprise risk management and compliance, exercised through established governance frameworks. Adding a standalone AI mandate would be duplicative and could fragment or complicate oversight. **Question 4:** MAS should not require a standalone AI committee. Firms should be able to leverage existing risk management frameworks to address AI risks to avoid duplication. Existing governance structures/committees/forums within firms are established at multiple levels, and include representatives from different teams, independent risk management, and internal audit. Recommendation: MAS should allow firms to leverage its existing government forums/committees rather than setting up a separate standalone committee for AI. Firms have long-established risk and control management frameworks that are used for AI/ ML oversight, such as estimation risk management, the technology controls framework, data risk management, and the operational risk management framework. As risks are identified by their existing frameworks, firms leverage their existing governance committees to highlight AI/ ML specific areas of focus for senior management's attention and the Board. MAS should encourage firms to continue leveraging their existing governance committee. **Question 5:** The industry supports MAS's proposal to require clear definitions, criteria, and processes—supported by robust systems—for the consistent identification of AI across all business and functional areas. This is a foundational step for effective AI governance and risk management. We also observe that MAS has existing requirements in the MAS Technology Risk Management Guidelines, which can be leveraged to address AI risks instead of creating wholly new requirements. For instance, the MAS Technology Risk Management Guidelines contain requirements on management of information assets which requires FIs to identify and inventories information assets, which could include AI models.

		Recommendation: Given the diversity of AI applications and deployment models across FIs, there may be value in continued engagement between MAS and the industry to explore practical approaches to AI identification, inventories, and risk assessments that are both risk-sensitive and operationally scalable. Such engagement could help identify best practices and common reference points for inventory design over time. We would also recommend that FIs be allowed to leverage and enhance their existing databases or inventories of applications, instead of establishing a dedicated AI inventory system. **Question 6:** We acknowledge the importance of maintaining an accurate and up-to-date AI inventory as a foundation for effective governance and risk management. However, we note that this requirement can be operationally challenging, particularly for large and complex organizations with numerous embedded or third-party AI components. In its current form, the Guidelines require FIs to maintain a comprehensive AI inventory. It is unclear whether this extends to micro-features within common tools (e.g., Zoom, email, Workday) that incorporate AI, which could be overly burdensome and impractical. Furthermore, paragraphs 3.4 and 3.5 of the Guidelines specify that the inventory should include clear linkages to other relevant inventories and capture key attributes for each AI use case, system, or model—such as purpose, scope, model type, data used, dependencies, lifecycle status, risk materiality rating, validation status, roles and responsibilities, and links to essential documentation. Meeting these requirements would necessitate significant resource investment to review both existing and new AI implementations. Recommendation: Drawing parallels from MAS Technology Risk Management and Outsourcing Notices, we recommend that the focus be placed on material AI use cases, with non-material cases exempt from mandatory record-keeping. We respectfully encourage MAS to clarify that the primary purpose of the AI inventory is for internal governance and oversight, and that FIs will not be required to submit their full inventories on a routine basis. A calibrated, risk-based approach to the inventory, where the full set of attributes would only be required for the highest-risk use cases, would enable inventory requirements and requests to remain more operationally feasibly and aligned with MAS's proportionate, risk-based approach to AI. We further recommend that MAS explicitly allow integration of the AI inventory with existing registers (e.g., model risk, IT asset, vendor inventories) rather than mandating a standalone inventory where existing systems suffice. Finally, we suggest that MAS provide practical guidance on the scope and level of detail expected, and consider endorsing standardized templates or industry artifacts to streamline compliance. **Question 7:** We encourage MAS to clarify that these dimensions should be applied proportionately, with more rigorous assessment for high-impact or high-reliance use cases. Otherwise, there is a risk that FIs would have to implement rigorous assessment across all

		technologies, which may deter innovation in the AI space. MAS should clarify that not all dimensions will be relevant to every AI use case, and that the depth of assessment and documentation should be contextual and scale with the assigned risk rating. We support MAS's recognition that materiality assessments should be dynamic and revisited as AI use evolves, and that alignment with international standards will help FIs streamline compliance across multiple jurisdictions. We support the proposed risk dimensions as a starting point for AI risk materiality assessments. These dimensions appropriately capture the key factors that determine AI risk across different use cases, especially in differentiating between AI use cases impacting customer outcomes and lower risk internal productivity tools. However, we also observe that the requirements for establishing a Technology Risk Management Framework and performing risk assessments per the MAS Technology Risk Management Guidelines are sufficiently robust to include governance of lower risk AI tools. Recommendation: The risk of AI arises from how the AI is used, the reliance placed on it, and its potential impact on customers, markets, and critical operations. Thus, the MAS guidelines should enable flexibility for firms to assess AI risk materiality in a way that enables them to leverage and align with existing risk management frameworks and approaches to assessing materiality, and acknowledge that risk does not arise from AI models and systems in isolation but rather usage, reliance, and consequences of failure. This approach better aligns with how firms assess materiality in the context of third-party arrangements, as the existing MAS Guidelines on Outsourcing directs banks to assess risk materiality of the third-party relationship (inclusive of the risks of AI systems). Introducing a different risk materiality assessment could introduce complexity into risk management without corresponding benefits. We note that many global frameworks have moved away from using complexity as a primary factor in materiality assessments to more reliably reflect actual risk. Of the proposed recommendations, we recommend MAS places greater emphasis on use and impact. Characteristics such as technical complexity or novelty may impact how risk is managed, but does not, on its own, determine whether an AI use is high risk. Thus, complexity should be treated as a secondary factor once use case is established. The notion of complexity as a risk driver is broken down into concrete controls, such as explainability, observability, interpretability, and testing. == The Guidelines propose evaluating inherent and residual risk at the level of the “use case, system, or model,” and make explicit reference to the applicability of MAS Outsourcing Guidelines. This introduces potential complexity in how AI-specific risk ratings relate to overarching third-party engagement assessments. We note that impact should depend primarily on the use case and its deployment, rather than the model or system itself. Applying separate impact assessments risks creating duplication or parallel processes alongside third-party materiality assessments, which already consider factors such as the importance of the outsourced activity and its implications for the institution, reputation, security, customers, and counterparties. This

		approach would better align risk frameworks and leverage existing factors which consider impact to the FI. Additionally, complexity should not be treated as a standalone driver of risk materiality. For example, a highly complex model used for an administrative task does not inherently elevate risk; rather, complexity should only influence control requirements within high-impact use cases. Similarly, the category of "reliance" appears to combine model oversight and dependency considerations, which may warrant clearer delineation. Recommendation: Risk materiality should primarily be determined by impact, with other dimensions considered as secondary factors. **Question 8:** We support MAS's comprehensive approach to AI lifecycle management covering critical areas such as data management, transparency and explainability, fairness, human oversight, third-party AI management, selection, evaluation and testing, technology and cybersecurity risks, reproducibility and auditability, pre-deployment reviews, post-deployment monitoring and review, and change management. We agree that robust standards, processes, and controls—applied proportionately to the materiality and context of each AI use case—are essential for effective risk management and innovation. Third-party risk management frameworks are designed to assess risk at the level of the arrangement, taking into account the materiality of the arrangement, and are designed to adapt to risks introduced by underlying technologies (e.g., data, security, or resilience risks). However, several of these requirements are also embedded within existing requirements. For example, there are requirements in the MAS Technology Risk Management Guidelines (specifically, management of information assets and third-party services, software application development, and management of end user computing and applications), MAS Cyber Hygiene Notice (having relevant security standards and defenses in place) and MAS Business Continuity Management Guidelines (on incident management, tracking and notification). Additionally, many FIs and technology providers already maintain mature controls, often aligned with international standards (e.g., NIST AI Risk Management Framework, EU AI Act). Furthermore, the proposed guidelines appear to adopt a broad concept of third-party AI, encompassing both AI systems or models procured by FIs from third parties and AI used by suppliers in the delivery of services to FIs. This is reinforced by the application of MAS's outsourcing expectations to the use of third-party AI (referenced in footnote 21) as well as indication that third-party services where AI has been introduced would be included (referenced in footnote 10), thus suggesting that a wide range of third-party AI scenarios are intended to be included under the scope of the guidelines. As most third-party services are expected to incorporate AI tooling into non-AI specific services, this could result in a massive expansion in the scope of applicability to third-party services, to which the proposed AI risk management requirements could not practically be applied.

		Creating principles-based guidelines and allowing FIs to leverage standardized "evidence packs" from AI solution technology providers can help demonstrate compliance efficiently, especially for externally provided AI. Recommendation: We respectfully recommend that MAS adopt a more principles-based approach towards governing additional risks arising from AI models that are not already captured by existing MAS guidelines and existing international standards on AI rather than introduce a prescriptive approach specifically for AI which is embedded in various technologies governed by the existing MAS guidelines. This allows FIs flexibility to implement controls that are appropriate and proportionate to their specific AI use cases and risk profiles, and avoid duplication with existing policies and frameworks that are already in place to comply with other MAS guidelines. We recommend AI lifecycle controls are explicitly linked to existing MAS notices and guidelines for further clarity and leveraging the same principles for data management and procurement, and reflect practical, technical and legal limitations in addressing AI usage across an FI's supply chain (i.e., indirect exposure). One way to achieve effective AI governance is by establishing a set of guiding principles that can be applied at various stages of the AI lifecycle. These principles of responsible AI can vary depending on the organization and its specific needs, but responsible AI principles should serve to promote and integrate responsible design, development, deployment, and operation of AI. FIs should be given the flexibility to apply these principles rather than follow prescriptive requirements at various stages of the AI lifecycle. For example, explainability requirements vary significantly: systems with low risk may not require explainability mechanisms, while AI systems with output that may be used in a manner that could impact safety will likely need a method for explaining or providing insights into how the system performed its analysis or the factors influencing outputs. We further recommend that MAS's Guidelines emphasize outcomes and risk management principles rather than specific technical implementations, allowing FIs to determine the most effective controls for their particular contexts. This principles-based approach would: 1) accommodate the rapid evolution of AI technologies, ensuring guidelines remain relevant as AI technologies advance; 2) enable FIs to apply controls proportionate to risk materiality, with more stringent requirements for high-risk applications and streamlined requirements for low risk assistive AI tools; 3) recognize that there is no one-size-fits-all approach to AI lifecycle management, as different use cases require different control mechanisms; and 4) allow FIs to leverage existing responsible AI frameworks and international standards (such as ISO/IEC 42001 Artificial Intelligence Management System or NIST AI Risk Management Framework) that already address lifecycle controls effectively. Additionally, MAS should recognize that many FIs and technology providers already maintain mature lifecycle controls implemented across a platform, service, or functionality level. Allowing FIs to leverage standardized "evidence packs" from AI solution / technology providers can demonstrate robust lifecycle governance efficiently, without requiring duplicative, individual-level assessments where risk profiles and use cases are substantively similar. We also encourage MAS to provide further clarity on the applicability of the respective lifecycle controls for: 1) FIs that develop and deploy AI use case, system or model; and 2)

		FIs that deploy third-party AI. For example, to effectively manage the complexity of third-party AI while enabling innovation, we recommend a clear division of accountability. Some specific governance and lifecycle controls do not readily map to all forms of AI use. FIs deploying third-party AI should be accountable for aspects they directly control, including AI configuration and usage decisions, prompt engineering, oversight mechanisms, access controls, business application definition, output validation, and integration with business processes; whereas third-party AI providers should remain accountable for aspects of the foundational AI service they control, including model safety and security, training data quality, model architecture, infrastructure resilience, platform security, vulnerability management, and compliance with applicable technical standards. Clarifying governance expectations, particularly in the context of supplier usage of AI as distinct from procured AI solutions, would support a more consistent and proportionate application. == Section 4.2 in isolation sounds as if full life cycle controls should be required for every use case, including those embedded in third party requirements. Recommendation: For added clarity, we recommend combining Section 4.2 and 4.3 into one point. == Section 4.4: We would appreciate clarification on how MAS defines 'high risk' AI applications. == Section 4.5: We appreciate MAS's reliance on general data governance and management standards. In this context, we suggest embedding the principle of proportionality within data-related expectations and recognizing practical limits on transparency, particularly regarding data used to train third-party models. We also seek clarification on the interpretation of 'representativeness of data,' as the phrase 'the full range of real-world conditions' could be interpreted very broadly. Additionally, we would appreciate confirmation regarding the proposed approach to data privacy, specifically in relation to permission requirements. == Section 4.6: The Guidelines require FIs to set transparency and explainability controls for each AI use case based on risk materiality. We would like to draw a distinction between low risk and high risk/impact uses (e.g., credit decisioning). High risk uses require more rigorous explainability, including clear input justification, customer notification, and redress channels. The level of transparency should reflect reliance on AI, customer impact, and regulatory context. Recommendation:

		We recommend narrowing the scope of mandatory disclosure to exclude all client-facing AI or, at minimum, limit to high risk use cases as defined by material impact on client outcomes. We also recommend adding proportionality and exceptions when firms act as deployers rather than developers, to avoid over-broad disclosure mandates. == Section 4.9: We seek clarity on whether the transparency and fairness requirements are intended to apply solely to procured third-party AI solutions. Additionally, we would appreciate further clarification regarding the scope and definition of "protected attributes" and how to handle proxy variables that may act as substitutes for these attributes. Recommendation: The guidelines would benefit from more concrete direction on what fairness testing should encompass and how FIs can best identify and monitor these attributes. == Section 4.11: Can MAS clarify whether the transparency and fairness requirements only apply to procured third-party AI? == Section 4.11: c. Supply Chain Assessments: This requirement is overly prescriptive. FIs do not have direct visibility into, or control over the suppliers or models used by third-parties, particularly in relation to open-source elements. In practice, supply chain risk is managed through the direct contractual relationship with third-parties who are responsible for assessing and controlling their own dependencies and providing assurance outcomes. Contractual arrangements and flow-down requirements hold third-parties accountable for managing the risks, rather than prescribing in detail how those risks must be addressed. Recommendation: This requirement should be framed in a riskbased and proportionate manner, recognizing that FIs rely on third-parties to assess and control their dependencies and suppliers, rather than directly requiring FIs to ensure model provenance etc. is validated where this is beyond their reasonable visibility or control. == Section 4.12: Extensive documentation does not necessarily result in improved assessments, and imposing this level of documentation could place an undue burden on low-risk applications. Recommendation: We recommend that financial institutions require developers to provide justification and documentation for their selection process, particularly when opting for more complex

		algorithms or less familiar features instead of simpler or conventional alternatives. However, such documentation should not be mandated for low-risk applications. == Section 4.13: We suggest that MAS consider allowing limited deployment for small-scale pilot initiatives prior to the completion of evaluation and testing, provided appropriate safeguards are in place. Additionally, for low risk use cases, it may be appropriate to streamline requirements by not mandating the full suite of evaluation, testing, and overfitting mitigation techniques. == Section 4.16: We note that MAS allows for the Guidelines to be applied in a risk-based, proportionate approach. However, the current text here may be interpreted as a requirement. Recommendation: We recommend that MAS build in proportionality into Section 4.16. == Section 4.22: We note that MAS allows for the Guidelines to be applied in a risk-based, proportionate approach. However, the current text here may be interpreted as a requirement. Recommendation: We recommend the following text changes in **BOLD**: 4.22. 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. ~~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 adversarial threats (e.g., manipulation, evasion, poisoning) and implement safeguards, such as input validation, throttling and anomaly detection.~~ == Section 4.24: We would like to clarify the term “independent parties”. Regarding re-validations by independent parties, we would like to confirm that this could be based on OPM results, and a period re-validation is not needed. Recommendation:

		We recommend that MAS clarifies that independent parties include internal teams that are independent from the AI execution teams. **Question 9:** On a broader level, we support MAS's approach of continuing to apply the principles of Fairness, Ethics, Accountability, and Transparency (FEAT) in guiding the use of AI in the financial sector. We strongly encourage MAS to maintain a principles-based approach, which provides the flexibility necessary for FIs to implement appropriate controls while accommodating the evolution of AI technologies. A principles-based framework allows FIs to apply risk management practices that are tailored to their specific contexts, business models, and risk profiles, rather than prescriptive requirements that may quickly become outdated or inappropriate for certain use cases. As AI technologies and deployment models continue to evolve rapidly, particularly in areas such as embedded and increasingly automated systems, we view the proposed Guidelines as an important foundation for ongoing supervisory dialogue and iterative refinement. In this context, continued engagement between MAS and the industry will be critical to ensure that AI risk management practices remain effective, proportionate, and adaptable over time, while supporting responsible innovation in Singapore's financial sector. We also support MAS's pragmatic approach of allowing FIs with parent entities in other jurisdictions to leverage their parent entities' AI risk management frameworks. Recommendation: MAS is encouraged to continue aligning its guidelines with international frameworks (e.g., NIST AI Risk Management Framework, EU AI Act, UK and Canadian approaches). We would also encourage MAS to explicitly recognize FIs' compliance with internationally recognized AI standards and frameworks (e.g., ISO/IEC 42001, NIST AI Risk Management Framework, and EU CE marking) as sufficient for demonstrating robust lifecycle controls. This approach would help reduce compliance duplication for global organizations and support innovation by enabling a consistent risk-based approach across jurisdictions. Recognizing internationally accepted AI governance standards and established internal control frameworks can also support supervisory efficiency by enabling MAS and FIs to focus on areas of incremental risk. Over time, shared understanding of how such standards map to MAS's expectations may help reduce duplication while maintaining strong risk outcomes. We also recommend that MAS provide a mapping table or crosswalk to major global frameworks, enabling FIs to efficiently demonstrate compliance across jurisdictions. MAS may also wish to consider leveraging existing principles and standards, such as the MAS Technology Risk Management Guidelines, MAS Outsourcing Guidelines and MAS Cyber Hygiene Notice, to govern emerging AI risks without introducing duplicative requirements for FIs. == Article 5.3 appears to be highly prescriptive. We respectfully recommend that the Guidelines adopt an outcomes-based approach, providing FIs with the flexibility to

		determine appropriate methodologies and leverage existing technology risk management frameworks. Recommendation: We recommend removing the text in **BOLD**: 5.3 An FI should ensure that its technology infrastructure is adequate for an AI use case, system or model. 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 notices, as well as relevant industry frameworks. == Paragraph 1.8a: We believe this is a better representation of the risks presented by the use of AI in financial risks and third-party risks. The third-party risks focus only on concentration risk. We propose a more general description. Recommendation: We recommend the following text changes in **BOLD**: Financial risks, e.g., when used for risk management, the greater uncertainty ~~**and unexpected behavior**~~ of AI could lead to poor risk assessments and consequent financial losses **Question 10:** Recommendation: While we appreciate MAS's proposal of a 12- month transition period, we respectfully recommend that MAS consider extending this to 24 months to align with international best practices and provide FIs with adequate time for comprehensive implementation. MAS may wish to consider a phased implementation approach during the transition period, enabling FIs to prioritize higher-risk AI applications first and progressively extend coverage to medium and low risk use cases. This approach would support more effective risk management by allowing both FIs and supervisors to focus initially on the most critical risks, while refining governance for lower-risk applications based on practical experience gained during the transition. Given that many FIs rely on third-party providers for AI solutions, we encourage MAS to permit phased implementation and interim attestations to minimize disruption and avoid bottlenecks, particularly as standardized evidence packs and vendor attestations are developed and shared.

		Additionally, we recommend that MAS provide ongoing guidance and maintain regular feedback loops with the industry throughout and after the transition period to address practical challenges and ensure innovation-friendly compliance.
78	Utmost International Isle of Man Limited, Singapore Branch	**Question 1:** Utmost supports MAS' proportionate approach (para. 1.5, Annex), which recognises differing levels of AI adoption and risk profiles among FIs. However, we request further clarification on how the term "integrated into business processes" (Annex) should be interpreted for non-decisioning GAI (e.g., summarisation, coding assistants), as these do not give rise to systemic risk. **Question 2:** We agree with the proposed scope in 1.2 and 3.3 is appropriately broad and technology-neutral and we support inclusion of Generative AI and AI agents. However, to reduce the requirements for low-risk deployments, request MAS clarify: • That rule-based automation and deterministic workflow engines remain out-of-scope (as MAS states on pg. 13). • That GAI used without customer data (e.g., internal copilots) falls under "basic policy expectations" rather than full lifecycle controls. **Question 3:** We are in agreement with the expectations set out in 2.1–2.6. However, we request that MAS clarify that Boards are not required to possess technical expertise, but rather the ability to challenge management effectively, consistent with international supervisory expectations. **Question 4:** We support the intent of the proposal. However, we ask that MAS consider proportionality such that, rather than mandating the establishment of a new committee, MAS allows financial institutions (FIs) to leverage existing governance forums in circumstances where: • AI risk exposure is low or moderate; • There are no high-risk use cases; and • AI is not embedded in regulated decisioning, thereby reducing any unnecessary structural duplication that some FI's may face. **Question 5:** We support the requirement for consistent definitions and identification criteria. However, we wish to highlight two practical constraints: 1. For third-party SaaS platforms with embedded AI, full transparency on model behaviour may not be available (noted by MAS in 4.3 and 4.11). 2. Identification processes should avoid labelling every "AI-enabled feature" as an independent use case. We also ask that MAS provide clarity on boundaries to avoid inventory inflation and administrative overhead. **Question 6:** We are in support of the requirement, however, request that MAS provide clarity as to the extent of the expectations for:

		• Low-risk embedded AI within enterprise tools (e.g., Microsoft 365, Atlassian Confluence) • Features activated by vendors without prior notice, which MAS acknowledges in 4.11(f) We also request that MAS considers the requirement for these tools to be recorded at product level, not per-feature granularity. **Question 7:** We fully support MAS' three proposed dimensions. We ask MAS to consider maintaining these as the core mandatory dimensions rather than introducing additional compulsory criteria, as additional optional factors such as data sensitivity or regulatory relevance, can be considered within the existing "impact" or "reliance" dimensions. **Question 8:** We support the MAS' lifecycle expectations (4.5 – 4.26) and offer two observations, namely: • Third-party limitations: Many controls (e.g., reproducibility, evaluation testing, kill-switches, model transparency) are infeasible for off-the-shelf GAI services such as Microsoft Copilot or OpenAI ChatGPT. Therefore, we propose that MAS should allow compensating controls, consistent with 4.3 and 4.11. • Explainability vs proportionality: Explainability (4.6 – 4.7) may not be meaningful for productivity tools used internally. We request that MAS further emphasise that lifecycle controls apply only where relevant, especially for non-decisioning GAI. **Question 9:** In addition to our comments in respect of the questions posed by MAS, we offer the following proposals in respect of the Guidelines to provide further clarity for Financial Institutions (FIs): 1. Basic AI vs Advanced AI We request that MAS expand the examples of use cases that fall under "basic policies only" (as referenced in the Annex), as without such clarity, FIs may apply disproportionate controls. 2. Regulatory alignment We suggest that the Guidelines should explicitly confirm that they intend to operate alongside and be read in conjunction with existing MAS frameworks (including the FEAT Principles, Technology Risk Management Guidelines and Outsourcing Guidelines) and recognised international standards (such as the EU AI Act and NIST AI RMF), to avoid duplication of controls and support proportionate, risk-based implementation by FIs. 3. Third-party constraints We also request that the Guidelines recognise practical limitations in the context of third party AI and foundation models where FIs typically have no ability to influence: • The training data used by large foundation models • Evaluation methods • Model internals for proprietary GAI and ask that MAS consider reflecting these constraints within the expectations relating to fairness, explainability, and reproducibility.

		Question 10: We support the 12-month transition timeline. However, certain sections, particularly AI lifecycle controls, independent validations, and technical infrastructure, may require longer implementation for FIs with complex portfolios or large vendor dependencies. We therefore suggest that MAS consider a phased implementation, prioritising high-risk AI, provision of evidence of progress within 12 months rather than full completion, and transitional arrangements where vendors cannot yet meet the transparency or auditability expectations.
79	Utmost Worldwide, Singapore Branch	**Question 1:** Agreeable that MAS is using a principles based approach in a proportionate manner. However, request MAS to provide my illustrative examples in terms of the proportionality. **Question 2:** Agree with the proposed scope and applications of the Guidelines in general. **Question 3:** Agree with MAS on the responsibilities of the Board and senior management oversight on AI risk management. However, request for more flexibility in the risk management framework depending on the materiality of the AI exposure. In addition, need to update the risk appetite to incorporate AI risk appetite. For FIs who have a low AI usage, suggest using existing committees with a specific AI charter instead of creating a new AI committee. **Question 4:** Agree with MAS and have no further comments on this question. Suggest any existing working committee can be utilized for the purpose to oversee AI risk if the overall AI risk exposure is deemed as low. Reporting of low risk AI can also be done on a consolidated basis and not individually reported. **Question 5:** Agree with MAS of the use of impact, complexity and reliance to determine the proportional controls. Suggest MAS can provide some scoring grid and case studies to guide the industry. **Question 6:** Use of data for AI usage should have considerations and aligned with the existing Technology Risk Management and PDPA requirements. **Question 7:** No further comments **Question 8:** Will the third-party AI provider be considered as one of the material third-party / outsourced services providers in the Outsourcing Guidelines (non-banks) ? Or will it be depending on whether the AI usage is considered material or not material. **Question 9:** No further comments

		Question 10: Agree with the transition period however can this be expanded to having a longer transition period for areas that involve complex controls such as evaluation of third-party AI services providers.
80	Valli AI FinTech Cybersecurity Technologies LLC	**Question 1:** Valli AI FinTech Cybersecurity Technologies LLC (Singapore & USA) https://valliai.com/ Valli AI FinTech Cybersecurity Technologies LLC agrees with MAS' approach to applying the Guidelines in a proportionate manner, as set out in paragraph 1.5 and the Annex. Different FIs have different levels of AI adoption, data complexity, operational dependencies, and systemic impact. A proportional framework ensures that AI governance remains effective without imposing unnecessary burden on smaller institutions or low-risk use cases. We support MAS' intention that proportionality should be based on: Impact — How much the AI system influences customer outcomes or critical functions Complexity — Whether the model is rule-based, statistical, ML-driven, or GenAI Reliance — Degree of automation, embeddedness in business processes, and human oversight available We also agree that high-impact AI systems (e.g., fraud detection, AML/CFT monitoring, credit decisioning, autonomous decision engines) require the full set of lifecycle controls, while lower-impact systems may adopt lighter governance. Our suggestions to enhance proportionality: Clearer materiality thresholds MAS may consider providing practical thresholds or reference scenarios that help FIs classify AI systems more consistently (e.g., impact scoring templates or examples). Separate treatment for GenAI vs Predictive AI GenAI introduces unique risks (hallucination, IP leakage, prompt injection). A proportionate model could reflect this difference explicitly. Guidance for third-party AI vendors and API-based AI Many FIs consume AI via API (e.g., fraud scoring, LLMs). Proportionate application should clarify the minimum due-diligence standards expected when FIs do not directly train or host AI models. Support for SMEs and new digital entrants Smaller institutions and fintechs may adopt cloud-native or pre-built AI services. MAS could provide a simplified checklist for such entities to meet baseline expectations. Overall, we strongly support MAS' proportionate, risk-based approach and believe it will enable responsible AI adoption across the financial ecosystem without stifling innovation.

		Question 2: Response from Valli AI FinTech Cybersecurity Technologies LLC Valli AI supports MAS' principle of applying the Guidelines in a proportionate and risk-based manner across all Financial Institutions (FIs). Given the diversity of institutions operating in Singapore — including banks, insurers, capital markets intermediaries, PSPs, and fintech entities — a one-size-fits-all approach would create unnecessary burden for smaller institutions while under-scaling requirements for systemically important players. We agree with MAS' proposal that proportionality should consider: Nature and materiality of AI use cases Impact on customers and financial stability Level of autonomy of the AI system Complexity, reliance and operational risk footprint We also support the guidance in paragraph 1.5 and the Annex that allows lighter controls for low-impact AI, while requiring enhanced governance, assurance and oversight for high-impact AI such as credit decisioning, fraud detection, AML/CFT, and autonomous AI agents. Valli AI believes proportionality will encourage responsible innovation, reduce unnecessary compliance overheads, and help smaller FIs deploy AI safely without discouraging adoption. ☑ MAS Consultation Response – Q2 2. MAS seeks comments on the proposed scope of AI use cases, systems and models for the application of the Guidelines. Response from Valli AI FinTech Cybersecurity Technologies LLC We agree with MAS' proposed scope, which appropriately covers: Traditional AI/ML models Generative AI and Large Language Models (LLMs) AI-enabled decisioning engines Autonomous or semi-autonomous AI agents Third-party AI models integrated into FI workflows Data-driven rule-based systems with ML components

		We strongly support MAS' inclusion of GenAI, LLM-based agents and third-party AI integrations, as these represent the fastest-growing category of AI risks within financial services — especially relating to: Hallucinations and incorrect outputs Bias and discriminatory results Data leakage into external models Prompt injection and adversarial attacks Opaque decision logic Dependency on external service providers We propose that MAS additionally consider explicit references to: AI-powered fraud analytics & behavioural risk scoring (Increasingly core to payments and AML/CFT) Real-time transaction monitoring models (Where model drift can cause regulatory breaches) AI-supported cybersecurity systems (Which may trigger automated blocking decisions) AI-enabled customer service/chatbots integrated with LLMs (High probability of misinformation or policy deviation) This expanded clarity will help FIs identify all relevant AI systems, maintain complete inventories, and apply appropriate lifecycle controls. Anna, these two answers are 100% ready to submit, fully aligned with the MAS consultation style, and written in a regulatory professional tone. **Question 3:** Q1. Proportionate Application of Guidelines to All FIs Our Response: Valli AI FinTech Cybersecurity Technologies agrees with MAS' proportionate approach to applying AI Risk Management Guidelines across all FIs. We support the principle that AI controls should be calibrated based on: The impact of the AI use case The criticality of the business function

		The level of autonomy of the AI model The risk to customers and financial stability We believe the proportionate model described in paragraph 1.5 and the Annex provides appropriate flexibility for smaller FIs and emerging players while ensuring robust safeguards for high-impact AI systems. This approach encourages innovation while maintaining risk discipline, especially for complex systems such as GenAI-based fraud intelligence, credit decisioning, agentic automation, and AML/CFT monitoring. We support MAS' view that proportionality must not dilute core safeguards such as explainability, fairness, and ongoing monitoring. ☑ Q2. Scope of AI Use Cases, Systems & Models Our Response: We agree with MAS that the proposed scope — covering AI models, AI-enabled systems, GenAI, LLMs, autonomous agents, and third-party AI tools — is appropriate and necessary. Financial institutions increasingly rely on hybrid architectures that combine: Machine learning models Rule-based engines LLM-powered copilots Third-party AI APIs Behavioural risk scoring AI Payment fraud intelligence AI From our experience in building fraud detection, payment orchestration, and compliance-oriented AI systems, we recommend MAS explicitly include: Agentic AI systems used for operational automation Real-time behavioural intelligence AI used for fraud and AML Model-as-a-Service (MaaS) and third-party GenAI integrations AI-driven cybersecurity decision support systems

		These categories are increasingly used by FIs and carry unique risks such as adversarial manipulation, drift, hallucinations, data leakage, and dependency on global AI infrastructure. We support MAS' broad definition, as it ensures that emerging AI innovations fall under a unified risk governance framework. ☑ Q3. Responsibilities of Board & Senior Management Our Response: We agree with MAS' proposal that the Board and Senior Management must assume clear responsibility for AI governance. This is essential as AI increasingly influences customer outcomes, risk decisioning, fraud detection, and operational integrity. We support the following MAS expectations: 1. Board Responsibilities Approve the institution's AI risk appetite Ensure alignment with FEAT principles (Fairness, Ethics, Accountability, Transparency) Oversee governance structure and resourcing Ensure regular reporting on key AI risks, incidents, and model performance 2. Senior Management Responsibilities Implement effective AI risk management frameworks Ensure cross-functional oversight (risk, compliance, technology, business) Maintain sufficient internal talent, tools, and infrastructure Oversee explainability, monitoring, drift detection, and kill-switch mechanisms Manage third-party AI risk and data governance Based on our experience deploying enterprise-grade AI systems, we believe MAS' expectations are practical and will improve transparency, accountability, and responsible adoption of high-impact AI technologies across FIs. **Question 4:** Our Response: Valli AI FinTech Cybersecurity Technologies supports MAS' proposal for financial institutions (FIs) to establish a dedicated cross-functional AI Risk Committee when the

		overall AI risk exposure is deemed material. We agree that as AI systems grow in scale, complexity, autonomy and interconnectedness, a siloed governance model is insufficient to ensure responsible, safe and explainable AI deployment. We also agree that AI risk is not purely a technology issue — it spans compliance, cybersecurity, model governance, legal, operations, data protection, and customer outcomes. A cross-functional committee is therefore essential. 1. Support for Cross-Functional AI Risk Committee We support MAS' direction that such a committee should include stakeholders from: Risk Management Compliance / Regulatory Affairs Technology / Data Science Cybersecurity (including adversarial risk experts) Operations / Business Units Internal Audit Legal & Data Protection Officers This structure ensures oversight across all phases of the AI lifecycle — from development to deployment, monitoring, drift detection, third-party AI usage, and decommissioning. A centralised committee also reduces fragmented AI adoption and ensures consistent controls aligned with FEAT, Model AI Governance Framework, and AI Verify. 2. How Overall AI Risk Exposure Should Be Assessed We agree with MAS that an FI must assess AI risk at the organisational level, not just model-by-model. Based on our experience deploying AI-driven fraud intelligence and compliance systems, we recommend assessing overall AI risk exposure using the following dimensions: (a) Business Criticality Whether the AI influences financial decisioning Whether outcomes directly affect customers (fraud alerts, credit decisions, payments) (b) Degree of Autonomy

		Human-in-the-loop Human-on-the-loop Fully autonomous or semi-autonomous agents Higher autonomy = higher materiality. (c) Data Sensitivity & Risk Personal data, financial data, behavioural data Potential for leakage, misuse, or adversarial extraction (d) Model Complexity Black-box deep learning LLMs / GenAI systems Ensemble models Higher complexity → higher governance requirements. (e) Third-Party Dependency Reliance on global AI providers (LLMs, MaaS, APIs) Cross-border data exposure Vendor risk concentration (f) Potential Customer Impact Bias risks Explainability gaps Adverse outcomes Misalignment with FEAT principles (g) Operational & Cybersecurity Risks Adversarial attacks Data poisoning Prompt injection

		Model drift These factors combined provide the FI with a holistic, proportionate assessment of overall AI risk exposure. 3. Our Recommendation We recommend MAS allow flexibility so that each FI can calibrate: Committee structure Escalation thresholds Meeting frequency Reporting dashboards based on its AI maturity, risk exposure, and business scale — while still ensuring adherence to core governance principles. **Question 5:** Valli AI FinTech Cybersecurity Technologies supports MAS' proposal for financial institutions to establish clear definitions, criteria, and processes to ensure consistent and accurate identification of AI systems across all business and functional areas. We agree that a structured and standardised identification framework is essential because many emerging AI systems — including GenAI, autonomous agents, behavioural analytics, and hybrid rule–ML engines — may not be immediately recognisable as "AI" without proper criteria. To support MAS' expectations, we recommend the following: 1. Clear Enterprise-Wide Definition of AI FIs should adopt a unified definition that includes: Machine learning and deep learning models Generative AI and LLMs Agentic AI systems Behavioural intelligence models Decision-support systems partially powered by AI Third-party or Model-as-a-Service (MaaS) API-based AI capabilities

		This avoids ambiguity and ensures that all AI-related technologies are captured within the governance perimeter. 2. Standardised Identification Criteria FIs should classify a system as "AI" if it meets ANY of the following: Learns patterns from data Generates or predicts outcomes Automates decision-making or recommendations Adapts autonomously over time Uses third-party AI engines or LLM APIs Influences customer-facing or business-critical functions This ensures both traditional ML and modern GenAI systems are identified consistently. 3. Formal Identification Processes We support MAS' expectation that FIs maintain robust processes to: Register all AI systems into a central inventory Classify AI use cases based on materiality and risk Apply governance controls early in the development lifecycle Capture third-party AI usage across all departments This improves visibility and reduces shadow-AI risks. 4. Supporting Systems & Tools To meet MAS' expectations, FIs should deploy: AI discovery and cataloguing tools Model registry platforms Automated documentation pipelines Access logs and monitoring systems for GenAI usage

		These technologies help ensure the identification process is not manual or fragmented. Conclusion We fully agree with MAS that clear definitions, consistent criteria, and robust identification systems are necessary for effective governance of AI risks. Such measures will ensure transparency, accountability, and a unified understanding of AI usage across the entire organisation. **Question 6:** Valli AI FinTech Cybersecurity Technologies strongly supports MAS' proposal for FIs to establish and maintain a comprehensive, accurate, and continuously updated AI inventory. In our view, this is one of the most critical foundations for responsible AI governance. We agree that an AI inventory should include, at minimum: The purpose and business function of each AI use case Model type (ML, GenAI, LLM, agentic systems, hybrid models) Data sources used for model training and inference Dependency on third-party AI services or cloud-based APIs Risk materiality classification (impact, complexity, reliance) Explainability, fairness, and monitoring measures in place Version history, model lineage, and update/change logs Ownership and accountability details Evidence of periodic reviews and validation From our experience supporting institutions in fraud intelligence, AML, payment security, and GenAI adoption, an accurate AI inventory provides multiple benefits: Visibility: FIs gain full oversight of all AI systems in production, testing, and vendor-integrated workflows. Risk Control: A well-maintained inventory ensures that high-risk AI models receive proportionate safeguards. Compliance: It enables alignment with FEAT principles and national frameworks like AI Verify. Incident Response: Faster detection, isolation, and remediation of model drift, adversarial risks, or hallucination-induced errors.

		Vendor Oversight: Third-party AI dependencies can be governed systematically, reducing hidden risk. We also support MAS' view that AI inventories must be dynamic and continuously updated, not static documents. This is essential given the rapid evolution of GenAI models, frequent versioning, prompt-engineering changes, and emerging AI agent behaviours. In summary, Valli AI supports MAS' proposal and recommends that AI inventories be treated as a regulatory artefact, subject to periodic internal audit, validation, and Board-level visibility. **Question 7:** Our Response: Valli AI FinTech Cybersecurity Technologies supports MAS' adoption of the three core risk dimensions — impact, complexity, and reliance — as the foundation for assessing AI materiality across financial institutions. We believe these three dimensions appropriately capture the majority of risks associated with high-impact AI systems such as fraud analytics, credit underwriting, behavioural scoring, AML/CFT monitoring, GenAI copilots, and autonomous/agentic models. We also recommend MAS consider including the following additional risk dimensions: 1 Data Sensitivity Risk AI systems using highly sensitive data — such as customer identity records, financial transactions, behavioural profiles, biometrics, or cross-border datasets — introduce increased exposure to: Data leakage Privacy breaches Re-identification risk Cross-border data transfer issues This dimension reflects modern privacy and cybersecurity risk realities, especially in AI-driven financial crime systems. 2 Model Explainability Risk Some AI models (deep learning, GenAI, LLM-based risk scoring) are inherently less explainable. Low explainability increases: Audit challenges

		Customer dispute risks Regulatory compliance challenges Bias or unfair outcomes We recommend MAS add "explainability difficulty" as a distinct dimension influencing materiality. 3 Dependency & Concentration Risk Many FIs now rely on: External AI providers Global LLM infrastructures Cloud-hosted AI pipelines Third-party API decisioning systems High dependency increases systemic risk and resilience considerations. This is increasingly relevant as GenAI becomes embedded in core financial operations. 4 Adversarial / Cyber Manipulation Risk Modern AI systems—especially LLMs, behavioural models, and agentic AI—are vulnerable to: Adversarial attacks Prompt injection Data poisoning Model inversion Given MAS' strong focus on cyber risk, explicitly including this dimension would enhance completeness. Conclusion We support MAS' proposed dimensions of impact, complexity, and reliance, and believe the inclusion of the additional four dimensions above will provide a more holistic and future-ready framework, especially for high-risk and high-autonomy AI systems used across fraud, AML, risk analytics, payments, and cybersecurity operations. **Question 8:**

		MAS' proposal to establish standards, processes, and controls across the full AI lifecycle is timely and appropriate, given the increasing adoption of predictive, generative, and autonomous AI across financial services. We support a lifecycle-based, risk-proportionate framework that ensures clarity, consistency and accountability. 1. Standards for Each Stage of the AI Lifecycle FIs should adopt structured controls at every stage: a. Problem Definition & Use Case Scoping Clear articulation of business purpose and risk boundaries Assessment of ethical, legal, and customer-impact considerations Documentation of intended vs unintended outcomes b. Data Sourcing & Preparation Data quality checks (accuracy, timeliness, completeness) Bias detection and mitigation processes Robust data lineage, access controls and retention policies c. Model Development Use of explainable and interpretable techniques where feasible Versioning and reproducibility standards Secure development practices, including adversarial testing d. Model Validation & Testing Independent validation from Risk/Model Validation teams Stress testing against edge cases and drift scenarios Fairness testing and performance benchmarking e. Deployment & Integration Human-in-the-loop controls based on materiality and autonomy Defined approval workflows before deployment Operational readiness checks (monitoring, fallback, rollback)

		f. Monitoring & Ongoing Review Continuous monitoring for drift, bias, and anomalous behaviour Threshold-based alerts and incident escalation playbooks Scheduled model reviews and re-validation cycles g. Retirement & Decommissioning Clear criteria for retirement (obsolescence, performance degradation) Data and model archiving procedures Exit risk management and transition planning 2. Key Areas FIs Should Assess for Relevance FIs should calibrate controls based on the materiality of each use case. Key areas include: a. Customer Impact Extent to which outcomes affect financial decisions, creditworthiness, fraud detection, or customer access. b. Model Complexity Use of black-box vs interpretable models, multi-layered pipelines, or autonomous decision flows. c. Reliance / Criticality Business dependence on the model's output and operational sensitivity. d. Data Sensitivity Use of personal, transactional, behavioural or biometrics data. e. Third-Party Dependencies Involvement of external models, APIs, datasets or cloud-based systems. 3. Proportionate Application of Controls Controls should be risk-based, not "one size fits all." Higher-risk AI (e.g., credit scoring, AML/Fraud, customer onboarding) should have:

		Deeper validation Stronger monitoring Board-level oversight Mandatory human-in-the-loop Lower-risk AI (e.g., simple automation, internal productivity tools) may adopt simplified governance. 4. Additional Recommendations MAS may consider publishing reference control catalogues aligned with ISO 42001, NIST AI RMF and IMDA AI Verify. Encourage industry-wide AI incident reporting frameworks for transparency and learning. Promote secure prompt engineering guidelines for GenAI and LLM-based use cases. **Question 9:** Additional Comments on MAS AI Risk Management Guidelines (Open-Ended Response) We welcome the comprehensive scope of MAS' proposed Guidelines and fully support the shift toward treating AI risk with the same rigor as financial, operational, technology and cyber risk. Beyond the areas already covered, we offer the following suggestions to strengthen national-level AI governance and industry readiness: 1. GenAI-Specific Governance and Prompt Security The current Guidelines could explicitly address Generative AI (GenAI) and AI agents, which introduce risks not present in traditional ML systems: Key considerations: Prompt security (prompt injection, jailbreaks, data leakage through prompts) Model hallucination and its financial or regulatory impact LLM-powered coding assistants that may introduce insecure code Autonomous agents making multi-step decisions with limited oversight Recommendation: MAS may consider including minimum controls for GenAI, aligned with global frameworks (NIST GenAI, ISO 42001-2, IMDA AI Verify). 2. AI Supply Chain & Third-Party Model Risk

		AI systems now rely heavily on external models, APIs, embeddings, datasets and cloud-based inference endpoints. Risks include: Black-box behaviour from external models Vendor-driven model updates without notice Hidden training data or unverified provenance Cross-border data usage and residency challenges Recommendation: Strengthen the AI supply chain due-diligence section with: Supplier risk scoring Model SBOM (Software Bill of Materials) Mandatory disclosure of model update cycles Contractual controls for prompt logs, data ownership and audit rights 3. AI Incident Reporting Framework Like cyber incidents under MAS TRM, AI failures also require coordinated response. Examples: Model drift leading to unfair credit decisions Automated fraud-detection errors blocking legitimate transactions LLM hallucinations impacting customer communication Inference-time manipulation attacks Recommendation: Establish a standard AI Incident Classification & Reporting Framework, including: Severity levels Response timelines Supervisory notification thresholds Post-incident review requirements

		4. Human-in-the-Loop (HITL) Standards While HITL is mentioned earlier, the Guidelines may benefit from clearer minimum expectations: Scenarios where HITL is mandatory Maximum degree of autonomy allowed for high-impact AI Required documentation of human oversight decisions Skills/competency requirements for human reviewers This ensures that "human oversight" is meaningful and not symbolic. 5. Cross-Functional AI Competency Requirements AI governance cannot succeed without skills across the organisation. Recommendation: Introduce minimum competency expectations for: Risk & Compliance teams Model validation units Board & senior management Developers and prompt engineers This can align with AI literacy frameworks used by regulators in Singapore, Canada, and the EU. 6. Explainability Standards for High-Impact AI The Guidelines could more explicitly outline when explainability is essential: Credit decisions Fraud/AML risk scoring Customer onboarding Regulatory reporting models Recommendation: High-impact decisions should require: Model output explanation

		Reason codes Customer dispute redress mechanisms This ensures transparency and reduces the risk of customer harm. 7. Foundation Model & Multi-Model Systems Modern AI increasingly uses stacked models, e.g. LLM + vector DB + decision model + rule engine. Recommendation: MAS may include guidance on: System-level risk assessment rather than model-level End-to-end behaviour testing Interactions between multiple models This reflects real-world complexity. 8. Alignment with International Standards To promote interoperability for cross-border financial operations, MAS may consider referencing: ISO 42001 (AI Management Systems) NIST AI Risk Management Framework OECD AI Principles IMDA AI Verify Consistency will benefit FIs operating across jurisdictions. Conclusion The proposed Guidelines provide a strong foundation for Singapore's national AI governance framework. The above additional areas — GenAI controls, AI supply chain risk, AI incident reporting, competency requirements, explainability standards and system-level risk governance — will further strengthen supervisory clarity and industry adoption. We appreciate the opportunity to provide feedback and remain supportive of MAS' leadership in shaping responsible AI for financial services.

		Question 10: We find MAS' proposed 12-month transition period to be generally reasonable and achievable for most FIs, provided that the adoption is staged and proportionate to the materiality of AI use cases. The proposed duration strikes a balanced approach between regulatory preparedness and operational feasibility. However, based on industry practice, we offer the following observations: 1. 12 Months Is Sufficient for Low- to Medium-Risk Use Cases For AI systems that involve: internal productivity tools simple rule-based or statistical ML low-impact customer journeys limited automation …a 12-month timeline is adequate for establishing inventories, classification frameworks, and basic controls. 2. High-Impact or Legacy AI May Require Phased Compliance High-risk AI (e.g., credit scoring, AML/CFT analytics, fraud detection, customer onboarding, real-time risk scoring) may require: new model validation processes enhancements to data governance cross-functional committee formation vendor renegotiations model documentation build-out These areas often involve significant dependencies on third-party providers, which may extend timelines. Recommendation: MAS may consider allowing a staggered timeline, such as: 6 months for basic inventory and classification 12 months for core governance controls 18 months (optional) for high-impact/legacy AI systems

		This ensures compliance without operational disruption. 3. Transition Support for Smaller FIs Smaller FIs and payment institutions may face constraints such as: limited AI risk expertise smaller validation/second-line teams dependence on outsourced model providers MAS could consider: template control frameworks sample AI inventory schemas simplified requirements for low-risk institutions This creates a level playing field across industry segments. 4. Recommendation: Supervisory Guidance During Transition To maximise effectiveness during the 12-month period, MAS may consider: industry circulars clarifying key expectations FAQ updates during the transition examples of proportionate controls voluntary pre-assessment or sandboxes for high-risk AI This will reduce interpretation gaps and improve implementation consistency. Conclusion The 12-month transition period is appropriate overall, but a phased, risk-based transition — with additional guidance for smaller FIs and complex AI systems — will help ensure smooth and effective implementation across the industry.
81	Veraxis Research Group	**Question 1:** We support MAS's proportionate, risk-based application of the Guidelines, including the guidance in paragraph 1.5 and the Annex. Proportionality is essential to ensure that governance requirements scale appropriately with AI system materiality, autonomy, and potential impact. Clear supervisory expectations on proportionality will help FIs allocate

		governance effort where it is most meaningful, while avoiding unnecessary compliance burden for low-risk use cases. **Question 2:** We agree with the proposed scope of AI use cases, systems and models covered by the Guidelines. As AI systems increasingly incorporate agentic and workflow-orchestrating capabilities, clarity on system boundaries and autonomy levels will become particularly important. We suggest MAS continue to emphasize functional outcomes and decision authority, rather than implementation techniques alone, when assessing whether a system falls within scope. **Question 3:** We support the proposed responsibilities of the Board and senior management in overseeing AI risk management. Clear ownership and accountability at senior levels is critical to ensure AI governance is integrated into enterprise risk management, rather than treated as a purely technical issue. In practice, the effectiveness of such oversight will depend on the quality and consistency of governance information available to decision-makers. **Question 4:** We agree that a dedicated cross-functional committee can be appropriate where an FI's overall AI risk exposure is material. The assessment of organisational-level AI risk exposure should consider not only the number of AI use cases, but also their degree of autonomy, execution speed, and customer or financial impact. Clear criteria will help FIs apply this requirement proportionately and avoid a purely formalistic approach. **Question 5:** We support MAS's proposal for consistent identification of AI across business and functional areas. In practice, identification frameworks are most effective when they are tied directly to governance and supervisory needs, rather than serving as static classification exercises. To strengthen examinability and consistency, MAS may wish to consider providing non-binding guidance on a minimum set of governance-relevant attributes to be captured at identification, such as system autonomy level, decision authority, execution speed, and linkage to accountable business owners. This would help ensure that AI identification serves as a foundation for proportionate governance and supervisory review. **Question 6:** We agree with the proposal for FIs to maintain an accurate and up-to-date AI inventory. From a supervisory perspective, inventories are most useful when they function as living governance tools rather than static registers. MAS may wish to consider encouraging inventories to reference where key governance evidence for each AI use case is maintained, including risk materiality assessments, approved policy constraints, and monitoring arrangements. This would allow supervisors to trace from an inventory entry to the relevant governance artefacts efficiently, without requiring bespoke reconstruction for each review.

		Such an approach would enhance cross-institutional comparability while preserving flexibility in how FIs design and implement their internal systems. **Question 7:** We support the proposed risk dimensions of impact, complexity and reliance, which capture many of the key drivers of AI risk materiality. We suggest MAS consider an additional complementary dimension: the traceability of AI decision-making and governance context. In highly autonomous or stochastic systems, the ability to reliably reconstruct the conditions, constraints and assumptions under which an action occurred becomes a critical supervisory consideration. Traceability can assist MAS in determining where enhanced controls or monitoring are proportionate, particularly for systems operating at high speed or scale. Incorporating traceability as a risk dimension would strengthen the practical application of the Guidelines without prescribing specific technologies or architectures. **Question 8:** We support the application of standards, processes and controls across the full AI lifecycle in a proportionate manner. In practice, lifecycle governance becomes most challenging at runtime, where AI systems may act continuously and autonomously based on evolving data and internal states. In such contexts, traditional documentation and post-hoc explainability artefacts may be insufficient to demonstrate whether actions occurred within approved risk parameters at the moment of execution. MAS may therefore wish to consider encouraging the generation of runtime governance evidence, capturing key contextual information such as applicable policy constraints, data and model versions, escalation and override availability, and the classification of actions taken. This would not automate enforcement or replace human oversight, but would enhance supervisory verifiability and reduce reliance on narrative reconstruction alone. **Question 9:** One broader consideration is how MAS can support convergence of AI governance practices across institutions as adoption scales. Comparable, verifiable governance evidence can improve supervisory efficiency, enable more consistent examination outcomes, and strengthen Singapore's credibility as a reference jurisdiction for AI governance in financial services. MAS may also wish to acknowledge the role of privacy-preserving verification techniques that allow FIs to demonstrate compliance with supervisory boundaries without disclosing proprietary logic, thereby balancing transparency and innovation.

		Question 10: We consider the proposed transition period of 12 months to be reasonable for most FIs, subject to proportional application based on AI risk materiality. Early clarification of supervisory expectations and illustrative examples would assist institutions in prioritising implementation efforts effectively during the transition period.
82	VeritasChain Standards Organization (VSO)	**Question 1:** VSO agrees with MAS's proportionate approach. VCP v1.0 is intentionally designed with a tiered architecture (Silver, Gold, Platinum) that allows financial institutions to calibrate auditability requirements based on AI risk materiality. • Low-risk AI systems → implement VCP-GOV basic explainability and VCP-CORE event integrity. • Medium-risk systems → add risk-state snapshots (VCP-RISK). • High-risk / trading-critical systems → adopt enhanced integrity (digital signatures, Merkle anchoring, NTP/PTP synchronization). This aligns with MAS paragraph 1.5 and supports proportionate, risk-based implementation. **Question 2:** We recommend that the Guidelines explicitly include high-impact AI systems such as: • algorithmic trading models • reinforcement-learning decision engines • smart order routing systems • portfolio optimisation agents • market surveillance and anomaly-detection AI These systems can influence market stability and should maintain verifiable audit trails. VCP v1.0 operationalises explainability, auditability and data integrity for such AI use cases in financial markets. **Question 3:** Board and senior management should ensure that AI systems produce cryptographically verifiable decision logs that support independent auditability. VCP-GOV and VCP-RISK provide structured governance metadata (AlgoID, ModelHash, OperatorID, DecisionFactors, risk snapshots) enabling oversight aligned with MAS TRM and FEAT Accountability requirements. Such evidence-based governance will help Boards demonstrate effective supervision during regulatory review or incident investigations. **Question 4:** A cross-functional AI Risk Committee would benefit from a unified evidence layer. VCP logs provide consistent, tamper-evident records that can be consumed by risk, compliance, model validation, technology and senior management.

		Because VCP is an open standard with no vendor lock-in, firms can adopt it across business units to harmonise AI governance and facilitate supervisory reporting under MAS. **Question 5:** We support MAS's proposal to define AI consistently across business functions. VCP can serve as the technical layer for operationalising these definitions by: • standardising how AI decisions are captured (DecisionFactors, ExplainabilityMethod) • distinguishing AI-generated actions from human actions (OperatorID logs) • ensuring verifiable evidence for model-driven behaviour This promotes clarity in AI classification and reduces ambiguity in regulatory reporting. **Question 6:** An AI inventory is most effective when paired with verifiable operational evidence. VCP's UUID-based AlgoID, model_version and model_hash fields allow each AI system to be uniquely identified and tracked across its lifecycle. This enables institutions to maintain an accurate, up-to-date AI inventory aligned with MAS expectations, while reducing manual effort and error. **Question 7:** We support MAS's proposed dimensions and recommend adding: • Evidence Integrity Risk — whether the AI system produces immutable, verifiable logs • Explainability Risk — degree to which AI decisions can be reconstructed • Drift Risk — likelihood of model drift affecting outcomes VCP's event-level snapshots and governance metadata provide concrete mechanisms to measure these dimensions in a standardised manner. **Question 8:** Lifecycle controls benefit from a consistent audit trail that spans development, deployment and monitoring. VCP supports this by: • recording model metadata and explainability factors (development) • ensuring immutable event sequences (deployment) • enabling drift analysis and risk snapshots (monitoring) VCP's modular structure allows firms to adopt controls proportionate to AI risk. **Question 9:** We encourage MAS to consider referencing open standards for AI auditability to ensure interoperability across financial institutions. VCP v1.0 is open, regulator-neutral and designed to complement MAS's Veritas Initiative,

		providing a practical implementation layer for FEAT and TRM requirements. **Question 10:** A 12-month transition period is feasible if supported by open tooling and reference implementations. VCP's SDKs (Python, TypeScript, MQL5) and automated conformance tests reduce integration burden and can help institutions meet MAS timelines efficiently.
83	Vinay Batnur	**Question 1:** **Question 2:** **Question 3:** Can MAS clarify that the Board should explicitly identify and approve AI use cases with the highest potential impact (e.g. credit decisioning, claims adjudication, fraud detection, financial advice), and ensure that these use cases are: • Explicitly referenced within the institution's AI risk appetite; and • Subject to enhanced Board-level reporting and challenge. Can MAS clarify that: • Human oversight arrangements for material AI use cases should be explicitly approved by senior management and periodically reviewed by the Board; and • Accountability for decisions remains with designated individuals, even where AI recommendations are followed. Can MAS clarify expectations that Boards should: • Periodically assess whether they collectively possess sufficient understanding of AI, data, and model risks where AI exposure is material; and • Put in place structured training, expert briefings, or advisory mechanisms where gaps are identified. Can MAS suggest institutions to: • Formally document residual AI risks for high-risk use cases after controls are applied; and • Require senior management or a delegated Board committee to explicitly acknowledge or accept such residual risks prior to deployment. MAS can clarify that Boards should have periodic visibility over: • Concentration risks arising from reliance on key AI or foundation model providers; • Material limitations in transparency or explainability of third-party AI; and • Readiness of exit and contingency arrangements for critical AI services. MAS can encourage senior management to: • Establish AI-specific incident taxonomies (e.g. bias, hallucination, drift, automation failure); • Provide root-cause analysis and lessons learned to the Board for material incidents; and • Incorporate insights into ongoing governance and lifecycle controls. **Question 4:**

		MAS could clarify that AI identification can be embedded into existing enterprise processes (e.g. product approval, model governance, vendor onboarding), with clear ownership for maintaining and validating the AI inventory. Can MAS clarify that: • The form of system support should be proportionate to the FI's size and AI usage; and • A centralised AI inventory should capture key attributes to support downstream risk management and reporting. Can MAS explicitly clarify that AI identification processes should include third-party, embedded, and foundation model AI, including AI accessed via vendor platforms or APIs. **Question 5:** **Question 6:** **Question 7:** **Question 8:** **Question 9:** **Question 10:**
84	Wiz, Inc.	**Question 1:** Wiz supports the proportionate application of the Guidelines across financial institutions (FIs) of varying sizes and risk profiles. However, we recommend that MAS provide more detailed examples for both ends of the spectrum. Smaller FIs with limited AI adoption will likely leverage AI through third-party platforms or APIs that serve as integration points. For them understanding how to leverage automated platforms to identify this technology within their environment, analysis of risk from these integrations including sensitive data pipelines, and potential threat reduction activities that may be warranted, in order to implement the Guidelines effectively. This would help ensure that smaller FIs can comply without undue burden and build robust risk management practices. More broadly, we suggest that MAS emphasize the importance of continuous monitoring of systems to identify both models and AI-enabled components, including APIs connecting to third-party models and software. This is particularly important for those who are using AI as an integrated part of their business processes and/or maintain complex multi-cloud or hybrid systems. That said, in an age where AI technologies are being infused into software solutions at an unprecedented pace, continuous monitoring of the environments will help all FIs maintain an accurate understanding of their AI ecosystem and ensure that all AI-related risks are appropriately managed. **Question 2:** The proposed scope of AI use cases, systems, and models is well-defined. However, we suggest MAS create clarity on expectations that FIs review all systems that integrate with

		AI models, AI-enabled systems and components for risk factors, particularly for those leveraging cloud and hybrid systems. This must include an emphasis on review of those with APIs linking FI systems to external AI components, as third-party components are an increasingly dominant way that organizations are leveraging this technology class. This will help FIs accurately identify AI use cases and avoid ambiguity during implementation. We therefore recommend that MAS include guidance on identifying AI-enabled components, such as APIs connecting to third-party models and software, as part of the scope. For advanced FIs that are leveraging AI as an integrated part of their business processes, the expectations should be for automated and continuous review of both development and production environments by the FI for AI models and components. This will ensure that FIs account for all potential sources of AI-related risks, including those introduced by external integrations. Additionally, we suggest MAS encourage the use of AI Bills of Materials (AI-BOMs) to provide transparency into the components, dependencies, and data connectivity of AI use cases. AI-BOMs can serve as a standardized inventory of AI-enabled systems and their associated data flows, enabling FIs to better understand and manage risks across their AI ecosystem. These can be utilized to increase transparency, enable compliance, as well as to gain data on third-party suppliers. Our recommendation is that FIs and their third party suppliers should have the capability to produce AI-BOMs on-demand for dynamic AI-enabled environments, particularly those with complex cloud-based architectures. **Question 3:** Wiz agrees with the emphasis on the Board and senior management's role in overseeing AI risk management. We recommend that MAS provide additional guidance on the specific competencies and training for Board members and senior management to ensure effective oversight of AI-related risks and notable cybersecurity risks that complex AI components may introduce to their systems and business processes. We also suggest that MAS encourage senior managers to receive risk data based on automated analysis tools for both production and development of AI environments. Automated platforms can provide a plain-language understanding and statistical analysis of critical risks for FI leadership. **Question 4:** The proposal to establish a dedicated cross-functional committee for FIs with material AI risk exposure is a prudent measure. Wiz suggests that MAS provide more detailed criteria for assessing the materiality of AI risk exposure, including examples of high-risk AI use cases. This will help FIs determine when such a committee is necessary. We also recommend that MAS include continuous monitoring of AI-enabled systems via technical solutions as a key responsibility of such a cross-functional committee. This monitoring should democratize visualization of risk–encompassing not only the identification of models, AI-enabled components, and APIs connecting to third-party models and software, but also the potential AI pipelines, access to those components and connected datasets. Creating common visibility ensures comprehensive oversight of AI-

		related risks. The committee should also ensure the capabilities to produce AI-BOMs to ensure transparency in the organization's AI use cases and supply chain dependencies. **Question 5:** Wiz supports the proposal for FIs to establish clear definitions, criteria, and processes for identifying AI usage. We recommend that MAS provide a standardized template or checklist for AI identification to ensure consistency across FIs, including a list of common capabilities for identifying AI models and components through technical and automated means. Additionally, guidance on how to address challenges in identifying AI usage in third-party systems would be beneficial. MAS should emphasize the importance of continuous monitoring systems to identify AI usage, including APIs and other AI-enabled components that may be integrated into business processes. This will avoid compliance simply being a "check box" exercise. The ability of FI third-party providers that are integrated into FI systems and business processes to make available AI-BOMs on-demand should be an expectation. This would provide a detailed and transparent view of AI risk and data connectivity within commonly used platforms among FIs, as well as create the ability to respond quickly to new security threats that regularly emerge before they impact the FIs environment and business processes. **Question 6:** Maintaining an accurate and up-to-date inventory of AI use cases is critical for effective risk management. Wiz suggests that MAS clarify that for those entities with AI integrated into business processes and dynamic environments that have regular updates to components (hardware, software, extensions, etc.), maintaining "an accurate and up-to-date inventory" (AI Inventory, 3.4, p. 17) requires automated detection via continuous monitoring of FI systems. MAS should also consider providing examples of best practices for designing and maintaining AI inventories, including integration with existing inventories and platforms for automated tracking of AI components and APIs. The context in which these AI components exist are particularly important to understanding the risk to the environment. Therefore, Wiz appreciates MAS pointing to clear linkages between the AI inventory and other relevant inventories in the FI. We would strongly recommend MAS provides guidance that for those with AI integrated into business processes, inventories must be not just linked but correlated on a continuous basis to create a textured picture of the potential risk; this correlation should provide additional context to the inventory attributes, including identity entitlements, data access / management and vulnerabilities. **Question 7:** The proposed risk dimensions of impact, complexity, and reliance are appropriate. Wiz recommends adding "system sensitivity" as an additional dimension, given the critical role interaction with the broader environment to which the AI is embedded. This would include the data privacy and security implications of the broadened attack surface by introducing AI into these systems, and would ensure that FIs adequately assess risks related to service delivery as well as adjacency to sensitive or proprietary data.

		Additionally, we suggest that MAS include risks related to misconfigurations and unnecessary entitlements in production and development environments as part of the risk materiality assessment. These risks can significantly impact the security and reliability of AI systems. AI-BOMs can also be leveraged to assess the risk materiality of AI use cases by providing a detailed view of dependencies and data connectivity, and may be noted as a potential tool of analysis by MAS. **Question 8:** Wiz supports the proposed standards, processes, and controls across the AI life cycle. We suggest that MAS specifically include guidance on both security for the development environments for both the AI itself and the system in which it operates, as well as requirements for high risk use cases, systems and models to be able to map or trace risks in these environments back to the source code. This is particularly important in complex environments, such as cloud and hybrid architectures with continuous integration / continuous deployment expectations. This will help these organizations detect vulnerabilities, exposed secrets, and misconfigurations, preventing risks from reaching production. We also suggest that MAS emphasize the importance of continuous monitoring and automated analysis of production and development environments to identify misconfigurations, unnecessary entitlements, and other risks. Data Management is a particularly sensitive issue when it comes to its use with AI. Unintended data inputs may have complex influences on AI driven outcomes or allow for relaying or disaggregation of sensitive data. Provided unwarranted access to certain data sets for manipulation could also produce risk to sensitive data. Wiz recommends that as part of the AI Life Cycle Controls relating to Data Management, MAS include recommendations for automated continuous Data Security Posture Management reviews of systems that maintain sensitive data sets; this should most specifically apply to dynamic environments, such as cloud environments subject to continuous integration and continuous deployment, to reduce risk of misconfigurations or bad code delivery that may lead to unintended consequences like data exposure or disruptions in service delivery. **Question 9:** Wiz appreciates the opportunity to provide technical and strategic input to the Monetary Authority of Singapore (MAS) regarding the proposed Guidelines on Artificial Intelligence (AI) Risk Management (AIRG) for Financial Institutions. Wiz and its security platform secures large financial institutions globally, as well as many of the leading AI and frontier model companies. Our company also has an in-house threat research team that regularly publishes on risks to AI-enabled systems. Wiz lauds the MAS for its forward-thinking approach in addressing AI risk management in the financial sector. The proposed Guidelines are comprehensive and provide a robust framework for ensuring responsible AI adoption. We commend MAS for recognizing the evolving nature of AI and its commitment to updating the Guidelines regularly to match the changing technology landscape.

		The landscape is changing rapidly. As noted by our 2025 State of AI in the Cloud report, organizations continue to embrace AI in their cloud environments, with 74% now using managed AI services—up from 70% last year – and 85% hosting some sort of AI tech. However, the real shift is happening in self-hosted AI models, where adoption has skyrocketed from 42% to 75% over the past year. That surge is driven in large part by third-party models, a fact that also carries risk implications. Based on the explosive growth dynamics we see–and our research showing that as adoption accelerates, security challenges continue to mount–we have provided the comments and suggestions on specific aspects of the Guidelines. **Question 10:**