

Asia Securities Industry & Financial Markets Association
RESPONSE TO PUBLIC CONSULTATION
Proposed Advisory Guidelines on Use of Personal Data in Generative AI
(Issued by the Personal Data Protection Commission on 2 June 2026)

Cover Page

Particulars of the Organisation

Name of organisation	Asia Securities Industry & Financial Markets Association (“ASIFMA”)
Nature of organisation	Independent, regional trade association representing leading financial institutions from both the buy side and the sell side, including banks, asset managers, professional services firms and market infrastructure service providers
Submission prepared by	Rajah & Tann Singapore LLP, on behalf of ASIFMA
Registered / business address	12 Marina View, Asia Square Tower 2 Level 11, #11-01 Singapore 018961
Website	asifma.org

Particulars of the Contact Person

Name	Rishi Kapoor
Designation	Executive Director, Head of Technology and Operations
Organisation / firm	ASIFMA
Email address	rkapoor@asifma.org
Contact number	+65 8221 6370

A. Summary of main comments

Question 1 — Are there other examples of “digital barriers” under paragraph 3.4 that would be useful to clarify?

ASIFMA suggests that personal data acquired from commercial third-party data providers (e.g. prospect database providers) under licensing agreements is a further useful example of data behind a "digital barrier". ASIFMA requests the Commission to clarify how the notify-and-assess process in paragraphs 3.7 to 3.9 should apply where AI training data is obtained from a commercial vendor rather than scraped from a website. ASIFMA also requests further guidance on benchmarks for applying the "reasonable person" standard mentioned in paragraph 3.3. Further, ASIFMA seeks confirmation from the Commission that the notify-and-assess process only applies in situations where an organisation wishes to scrape personal data for the purpose of Generative AI model training, and that a collecting organisation may proceed with the scraping if the notified organisation fails to respond to the former's notification after a reasonable amount of time. Lastly, ASIFMA would appreciate the use of more illustrative examples to clarify how the notify-and-assess process should be carried out in practice.

Question 2 — Should organisations provide “AI-Specific Notifications” in situations other than Generative AI Model training and/or fine-tuning?

ASIFMA's primary submission is that the Commission should adopt a technology-neutral approach and avoid imposing a blanket "AI-Specific Notification" requirement. ASIFMA does not view all AI-related uses as inherently posing greater risks of harm to data subjects, considering that AI Models can be deployed across a wide range of use cases, and any potential concerns would already be adequately addressed through existing legal requirements and remedies. However, should the Commission nevertheless wish to retain the “AI-Specific Notification” requirement, ASIFMA submits that it should be confined to the initial training of Generative AI Models rather than extended to all AI-related uses including non-training operational uses. In addition, the Commission should permit general disclosures (e.g. that personal data will be used for “AI-related purposes” broadly) rather than prescribing an overly specific disclosure requirement and its associated content. Further, it should be clarified that the “AI Specific Notification” requirement, if implemented, would apply prospectively only, preferably through a grandfathering provision and, only where such grandfathering is not feasible, through a reasonable transition period for personal data already collected. ASIFMA also seeks confirmation that the “AI-Specific Notification” requirement may be met by the existing practices of banks and financial institutions which are already expected to provide certain AI-related disclosures under the Monetary Authority of Singapore's Fairness, Ethics, Accountability and Transparency (FEAT) principles.

Question 3 — Are there other best practices that would be useful to highlight at paragraphs 4.6 (substance of notifications) and 10.4 (complying with access and correction requests)?

ASIFMA has ongoing concerns about the feasibility of complying with access and correction requests in the Generative AI context. ASIFMA seeks confirmation that (a) organisations may adopt proportionate, risk-based or phased approaches in dealing with such requests; and (b) organisations acting reasonably will be treated as compliant even where full technical compliance is not presently attainable, and that compensating controls such as access restrictions, anonymisation or output safeguards are acceptable where deletion of erroneous data is not yet technically feasible in a Generative AI context.

Question 4 — Is there additional information on data protection safeguards that would be helpful for Model and System Providers to share with their downstream stakeholders?

ASIFMA emphasises the importance of having robust contractual frameworks which clearly allocate roles and responsibilities across Model Providers, System Providers and System Deployers, and

recommends that the Commission specify a baseline set of information that AI providers should minimally be expected to disclose to their downstream deployers. Besides having clearly allocated contractual responsibilities, ASIFMA submits that it would be important for organisations to document their internal practices clearly. This would facilitate the review of such practices at periodic intervals and allow organisations to enhance the effectiveness of their AI governance. ASIFMA also seeks clarification of the meaning of "data minimisation" in the context of Generative AI Models, and recommends that the Guidelines make clear that the roles of Model Provider, System Provider and System Deployer are not mutually exclusive, and that a single entity may potentially take on more than one of such roles at the same time.

Question 5 — Are there other agent-specific data challenges or risks that would be helpful to clarify at paragraph 9.5?

ASIFMA identifies several agent-related data protection risks but suggests that they be addressed in a technology-neutral, outcome-focused manner rather than by reference to whether a system is characterised as "agentic," given that the definition of AI agents remains unsettled. Further, ASIFMA would appreciate if the Commission can confirm that it would be sufficient for an organisation to engage in periodic assessments of its safeguards against such risks, and that continuous or real-time monitoring of the sufficiency of its safeguards would not be expected.

B. Detailed comments

Question 1: Are there other examples of "digital barriers" under paragraph 3.4 that would be useful to clarify?

1. One additional example of a digital barrier that would be useful to clarify is data that is obtained from commercial third-party data providers.
 - (a) Financial institutions frequently acquire personal data from third-party vendors (e.g. prospect database providers) under licensing agreements. Since such arrangements would typically involve the vendor's online databases being subject to both access controls and/or paywalls, they could be viewed as examples of "digital barriers".
 - (b) In this regard, it would assist members for the Commission to confirm if and how the notify-and-assess process set out in paragraphs 3.7 to 3.9 of the Proposed Guidelines should apply where the source of the data is a commercial vendor rather than a website which may be scraped.
 - (i) We understand that, pursuant to paragraphs 3.7 to 3.9 of the Proposed Guidelines, an organisation that wishes to scrape a website for data should first notify the website-owning organisation of its intention to do so. The website-owning organisation must then assess whether the data would constitute "publicly available data". If so, the website-owning organisation may disclose the data to the collecting organisation.
 - (ii) However, where the collecting organisation obtains its training data from a database maintained by a commercial third-party data provider, the data provider would already have knowledge of the collecting organisation's intention to use the data. The purposes for which the data may be used or disclosed would also have been set out in the parties' data licensing agreement.
 - (iii) Therefore, it would be helpful if the Commission could clarify and confirm:

- (A) that the notify-and-assess process should not apply where the parties have already entered into a licensing arrangement in respect of the data;
 - (B) that the terms of the licensing agreement would determine if the collecting organisation has discharged its Personal Data Protection Act 2012 (“PDPA”) obligations when collecting personal data in such scenarios; and
 - (C) whether the collecting organisation would be required to take any other steps to discharge its PDPA obligations when collecting personal data in such scenarios.
2. Relatedly, it would also be helpful for the Commission to provide further guidance on how the “reasonable person” standard in paragraph 3.3 is to be applied.
- (a) We understand that, pursuant to paragraph 3.3 of the Proposed Guidelines, personal data may only be used to develop a Generative AI Model if a “reasonable person” would consider the use of personal data appropriate in the circumstances. This is a reflection of the Purpose Limitation Obligation in section 18(a) of the PDPA.
 - (b) Considering the wide range of potential use cases, it would be beneficial to organisations if the Commission could provide some commonly accepted thresholds or benchmarks with reference to which “reasonableness” may be assessed when training Generative AI Models on publicly available data.
3. ASIFMA would appreciate confirmation from the Commission of the following points in relation to the notify-and-assess process set out in paragraphs 3.7 to 3.9 of the Proposed Guidelines:
- (a) that a collecting organisation may proceed with the scraping of personal data on a notified organisation’s site if the latter has failed to respond to the former’s notification after a reasonable amount of time; and
 - (b) that the notify-and-assess process only applies in situations where an organisation wishes to scrape personal data from a site for the purpose of training and/or fine-tuning a Generative AI Model, and not when data is intended to be scraped for other purposes.
4. ASIFMA would also welcome the use of more illustrative, worked examples (such as the one on page 9 of the Proposed Guidelines) to clarify how the notify-and-assess process¹ is to be carried out in practice, including examples which can provide further clarity on (a) sample notifications that can be used for the purposes of paragraph 3.7 of the Proposed Guidelines; and (b) how the digital barrier assessment factors found in paragraph 3.9 of the Proposed Guidelines may be applied. Given the widespread use of Generative AI in the financial sector, ASIFMA would also appreciate if such worked examples could feature use cases typically encountered by banks and financial institutions.

Question 2: Should organisations provide “AI-Specific Notifications” in situations other than Generative AI Model training and/or fine-tuning?

¹ Paragraphs 3.7 to 3.9 of the Proposed Guidelines.

5. ASIFMA's primary submission is that the Commission should avoid imposing a blanket requirement that all AI-related uses of personal data should be notified to data subjects through a special "AI-Specific Notification", simply because an AI System is involved.
- (a) A technology-neutral approach would be more appropriate. AI and Generative AI are merely extensions of existing data-processing technologies, and traditional systems can pose equal or greater risks to individuals if personal data processed by such systems is misused. ASIFMA does not view all AI-related uses as inherently posing greater risks of harm to data subjects, considering that AI Models can be deployed across a wide range of use cases, and any potential concerns would already be adequately addressed through existing legal requirements and remedies.
 - (b) As the "AI-Specific Notification" requirement in the Proposed Guidelines is a manifestation of the Notification Obligation under section 20(1) of the PDPA,² it is worthwhile to note the manner in which the Notification Obligation is ordinarily assessed under the PDPA.
 - (i) According to the Commission's Advisory Guidelines on Key Concepts in the PDPA,³ the PDPA does not provide for a specific form or manner in which an organisation is to discharge its Notification Obligation. Instead, organisations are allowed to determine the best way of doing so such that the individual is provided with the required information to understand the purposes for which his personal data is processed. The appropriateness of an organisation's form of notification depends on a range of factors, including:
 - (A) the circumstances and manner in which it will be collecting the personal data;
 - (B) the amount of personal data to be collected;
the frequency at which the personal data will be collected; and
 - (C) the channel through which the notification is provided (e.g. face-to-face or through a telephone conversation).
 - (ii) Therefore, it is a clear and established requirement that the manner in which the Notification Obligation should be discharged (including the required level of detail in the notification) should be calibrated to the context of the specific interaction in the circumstances.
 - (iii) As such, we submit that the Commission should avoid fixing or prescribing any particular standard or form for the notification of AI-related uses. Depending on the circumstances and the specific AI-related use, certain data subjects may require less detail about the way their personal data is being used, and others more. Imposing overly specific disclosures on all AI-related uses as a blanket rule therefore may not meaningfully enhance individuals' understanding of the associated risks, and may instead unnecessarily increase the administrative burden on organisations.
 - (iv) ASIFMA recommends that the Commission allow organisations to tailor their notifications according to their assessment of the risks associated with the specific use of personal data. Relevant factors may include, for example, whether the use is external-facing or purely internal, whether the data is processed transiently or retained for a period of time, whether the data is used solely for service delivery or also for product improvement or testing, and whether the use is aligned with the data subject's reasonable expectations.

² Paragraph 4.2 of the Proposed Guidelines.

³ Paragraphs 14.9 and 14.10 of the Advisory Guidelines on Key Concepts in the PDPA.

6. Should the Commission decide that an “AI-Specific Notification” requirement would still be necessary, ASIFMA submits that (1) the definition of “AI-Specific Notification” in paragraph 4.3 of the Proposed Guidelines should be limited to Generative AI Model development; and (2) the requirement for such AI-Specific Notifications should not extend beyond the use of personal data for the initial training of Generative AI Models. For operational uses of Generative AI Systems that do not involve training – such as inference, summarisation, retrieval, drafting, customer support, and screening for anti-money laundering purposes – a general notification of the purpose of processing should ordinarily suffice.
- (a) As the ambit of the Proposed Guidelines relates to the use of personal data in Generative AI (and not all forms of AI), the definition of “AI-Specific Notification” in paragraph 4.3 of the Proposed Guidelines should be limited to Generative AI Model development only, and the inclusion of the words “AI and/or” may be unnecessarily expansive and confusing for the purposes of the Proposed Guidelines. The Commission may wish to consider amending the terminology used in the Proposed Guidelines to “Generative AI-Specific Notification” to make the scope of such notifications even clearer.
 - (b) Generative AI Systems are deployed across a wide spectrum of uses with materially different risk profiles. The training of models can have broader and longer-term implications, including the persistence of personal data within a model and greater difficulty in its subsequent removal or correction.⁴ On the other hand, many non-training uses are contextual or transactional, may involve only transient processing, and present a different risk profile. Any requirement for a “specific” notification should be confined to higher-risk scenarios only, and the Commission should clarify the scenarios it regards as higher-risk for this purpose. This risk-based approach avoids imposing a disproportionate administrative burden on organisations.
 - (c) It would therefore be helpful for the Proposed Guidelines to draw a clearer distinction between: (i) the use of personal data to train a Generative AI Model; and (ii) the use of personal data solely to operate or support a Generative AI System without such data being used for model training (e.g. to perform inference, summarisation, retrieval etc), and to confirm that the notification approach applicable to each may differ depending on the circumstances. In this regard, it would also assist for the Commission to confirm that personal data collected by a System Deployer for AI processing (as opposed to training) purposes may be notified through a general notification.
 - (d) ASIFMA submits that any requirement to provide an “AI-Specific Notification” should be limited to the use of personal data in the initial training of Generative AI Models, where the impact on individuals is likely to be more significant. Extending the requirement to fine-tuning or subsequent refinement risks imposing additional complexity without a corresponding increase in meaningful transparency for individuals. This is especially as the fine-tuning process is an iterative one which may encompass a range of product / service improvement activities which may not be specific to Generative AI Systems, such as retrieval or ranking optimisation, testing and QA, and safety monitoring.
7. As to the content or degree of granularity required in AI-Specific Notifications, we submit that a high-level disclosure that personal data is used for an AI-related purpose should, in most cases, be sufficient. The Commission should avoid taking the approach of prescribing the content that should be included in such notifications, as the adequacy of the notification should be assessed

⁴ Paragraph 10.3 of the Proposed Guidelines.

on a case-by-case basis.

- (a) In today's economy, most organisations, including financial institutions, deploy AI across a wide variety of day-to-day operations (e.g. from general analytics purposes to assisting with the conduct of routine customer checks and procedures). Moreover, as AI continues to rapidly evolve, new use cases for AI emerge frequently which are quickly adopted by organisations. It would not be feasible, reasonable or beneficial for the AI-Specific Notification to include all the ways in which AI is used in an organisation's operations, or for a new AI-Specific Notification to be issued every time a new AI use case is adopted by an organisation.
 - (b) In addition, based on our members' practical experience interacting with customers, where an individual is averse to the use of their personal data for AI-related purposes, such individuals typically do not hold meaningfully distinct views between different AI-related purposes and their position is unlikely to change materially, even if they are informed of the specific manner in which their personal data would be processed by AI.
8. ASIFMA further seeks confirmation that any requirement to provide AI-Specific Notifications would apply prospectively only to personal data collected after the Guidelines take effect. In particular, we ask the Commission to confirm that organisations are not expected to issue new AI-Specific Notifications, or to re-obtain consent, in respect of personal data that has already been collected under previously compliant General Notifications and used for model training. Requiring organisations to revisit consent for data already collected would be disproportionate and operationally disruptive. Therefore, ASIFMA strongly recommends that the Commission include in the Proposed Guidelines a clear grandfathering provision, or, only where such grandfathering is not feasible, a reasonable transition period after further discussion with members.
 9. In terms of practical implementation of the AI-Specific Notification requirement, ASIFMA would also welcome the Commission's confirmation on whether a general update to an organisation's privacy notice reflecting its use of Generative AI would be sufficient, or if a separate notification to data subjects would be expected.
 10. To the extent that a separate notification is strictly required on every occasion, ASIFMA is of the view that this would be administratively burdensome given the rapid rate at which new AI use cases emerge and become adopted today. Providing notifications to data subjects too frequently would also contribute to data subjects' notification fatigue, reducing the overall effectiveness of notifications in helping data subjects provide informed and meaningful consent to their data processing. Moreover, ASIFMA notes that the practice of issuing separate AI-Specific Notifications does not seem to be aligned with that in other jurisdictions with mature privacy regimes, such as the EU, where AI-related uses may simply be disclosed to data subjects in the organisation's general privacy notices.
 11. In addition, ASIFMA would welcome confirmation from the Commission that the AI-Specific Notification requirement is substantively similar to the transparency requirements found in the Monetary Authority of Singapore's ("**MAS**") Principles to Promote Fairness, Ethics, Accountability and Transparency ("**FEAT**") in the Use of Artificial Intelligence and Data Analytics in Singapore's Financial Sector (or "**FEAT Principles**" in short).

- (a) Under the “transparency” section of the FEAT Principles (principles 12 to 14),⁵ financial institutions are expected to:
 - (i) proactively disclose the use of AI and data analytics (“**AIDA**”) to data subjects as part of general communications in order to increase public confidence;
 - (ii) provide data subjects (upon request) clear explanations on what data is used to make AIDA-driven decisions about the data subject, and how the data affects the decision; and
 - (iii) provide data subjects (upon request) clear explanations on the consequences that AIDA-driven decisions may have on them.
 - (b) Many banks and other financial institutions have already incorporated the FEAT Principles into their disclosure frameworks to customers.
 - (c) It would therefore assist if the Commission could clarify that such existing disclosure practices under sectoral regulations would be sufficient for the purposes of meeting the AI-Specific Notification requirement and associated consent and notification requirements found in the Proposed Guidelines.
12. Lastly, ASIFMA recommends that the Proposed Guidelines make clear that the statutory exceptions under the PDPA (such as the business improvement and research exceptions) remain fully applicable in the Generative AI context, and that if an organisation can avail itself of such exceptions on the facts, it would not be required to provide any purpose notification (AI-Specific or not) to its data subjects.⁶

Question 3: Are there other best practices that would be useful to highlight at paragraphs 4.6 (substance of notifications) and 10.4 (complying with access and correction requests)?

13. ASIFMA appreciates the Commission’s recognition of the many present-day challenges to facilitating Generative AI-related access and correction requests, including the issues brought about by the sheer volume of data used by Generative AI models and the current technical limitations in removing specific information from such models.⁷ However, we continue to have concerns on the feasibility of complying with access and correction requests in a Generative AI context which are not fully addressed by the Proposed Guidelines.
14. Under the PDPA, the Access and Correction Obligations lie with the organisation as data controller, and not with the data intermediary.⁸ In the generative AI context, System Deployers are often organisations (or data controllers)⁹ which rely on System or Model Providers as data intermediaries.¹⁰
15. While the Commission has acknowledged the difficulties faced by Generative AI System or Model Providers in complying with access and correction requests,¹¹ such risks would be amplified on the part of System Deployers, which face the direct burden of potentially breaching their Access and Correction Obligations if their System or Model Providers are unable (e.g. due to

⁵ See paragraph 8.2 of the MAS FEAT Principles.

⁶ See s 20(3)(b) PDPA.

⁷ Paragraph 10.3 of the Proposed Guidelines.

⁸ Paragraph 6.20 and 6.21 of the Advisory Guidelines on Key Concepts in the PDPA.

⁹ Paragraph 9.1 of the Proposed Guidelines.

¹⁰ Paragraphs 7.4 and 8.3 of the Proposed Guidelines.

¹¹ Paragraph 10.3 of the Proposed Guidelines.

technological or operational constraints) to facilitate the data subjects' requests. It would therefore be helpful for the Proposed Guidelines to expand further on how System Deployers can comply with their Access and Correction Obligations in the context of Generative AI.

- (a) In addition to the statutory exceptions to the Access and Correction Obligations mentioned at paragraph 10.2 of the Proposed Guidelines, other potentially relevant statutory exceptions to these Obligations under the PDPA would include:
 - (i) the data requested being opinion data kept solely for an evaluative purpose;¹²
 - (ii) the information which is the subject of the access request not existing or not being able to be found;¹³ and
 - (iii) the personal data which is the subject of the correction request being derived personal data.¹⁴

ASIFMA would welcome additional clarification in the Proposed Guidelines on how the statutory exceptions to the Access and Correction Obligations under the PDPA may apply in the Generative AI context.

- (b) It would also be helpful for the Proposed Guidelines to clarify how compliance with the Access and Correction Obligations would be assessed for an organisation that has acted in good faith in acceding to the data subject's request but is nonetheless constrained by the current limits of the technology.
 - (i) ASIFMA would welcome confirmation that, given the technical difficulties of complying with access and correction requests in the Generative AI context, organisations may adopt proportionate, risk-based or phased approaches to dealing with such requests, and that an organisation that acts reasonably in the circumstances would be regarded as compliant even where full technical compliance is not presently attainable.
 - (ii) For example, where an organisation has implemented upstream data accuracy controls, reviewed a request in good faith, and excluded the relevant erroneous personal data from future training runs, we propose that the Commission should take such facts into account when assessing the organisation's PDPA compliance, even if the organisation is unable (e.g. due to the present technical infeasibility of machine un-learning) to remove the erroneous data directly from an already-trained model. An organisation which adopts alternative measures (given present technological constraints) to meet its data protection obligations in good faith should be given due credit and safe harbour protections against any potential findings of non-compliance.
 - (iii) Providing express recognition of the above in the Proposed Guidelines would reward good practice by organisations without penalising organisations for limitations inherent in the present state of the technology.
 - (iv) We take this opportunity to seek clarification of the Commission's expectations regarding non-retention of personal data in the Generative AI context. In particular,

¹² Paragraph 1(a) of the Fifth Schedule and Sixth Schedule of the PDPA.

¹³ Paragraph 1(j)(iii) of the Fifth Schedule of the PDPA.

¹⁴ Paragraph 1(f) of the Sixth Schedule of the PDPA.

where data deletion is not technically possible or feasible, ASIFMA would welcome confirmation that the Commission would regard alternative or compensating controls – such as access restrictions or safeguards applied to model outputs – as appropriate risk mitigation measures.

- (c) ASIFMA welcomes the Commission’s acknowledgement of the technical challenges involved in facilitating access and correction in the Generative AI context, and the qualification that the relevant best practices apply “where reasonable”.¹⁵ Further guidance on how the “where reasonable” standard is to be applied would be beneficial, particularly where an organisation has limited ability to trace, retrieve or correct personal data once it has been processed within a Generative AI System.

Question 4: Is there additional information on data protection safeguards that would be helpful for Model and System Providers to share with their downstream stakeholders?

- 16. ASIFMA notes the importance of having robust contractual frameworks across Model Providers, System Providers and System Deployers. Given the layered nature of Generative AI ecosystems, the clear contractual allocation of responsibilities is an important means of avoiding ambiguity and managing downstream risk. Such contracts should clearly address:
 - (a) the allocation of roles and responsibilities between the parties, including whether each party is acting as an organisation or data intermediary in respect of each processing activity;
 - (b) the permitted uses of personal data, including whether such data may be used for model training, model improvement, testing, safety monitoring, or strictly for inference;
 - (c) data governance and security arrangements, including breach management responsibilities, retention and deletion requirements, audit rights, and the implementation of appropriate technical and organisational safeguards; and
 - (d) support for regulatory compliance, including cooperation in handling access and correction requests, withdrawal of consent, notification and consent responsibilities, and other obligations under the PDPA.

ASIFMA considers that encouraging parties to document these matters expressly would reduce ambiguity across the AI supply chain and strengthen accountability under the PDPA framework.

- 17. ASIFMA notes that both primary organisations and data intermediaries are subject to the Protection Obligation under section 24 of the PDPA. To enable System Deployers (typically data controllers) to discharge that obligation effectively, it would be of considerable assistance for the Commission to specify a baseline set of information that Model Providers and System Providers are minimally expected to disclose to downstream deployers.
 - (a) This is distinct from the approach of identifying certain categories of information that would be “good practice”¹⁶ or “helpful”¹⁷ to share with deployers.
 - (b) Such a minimum disclosure standard might encompass:
 - (i) the categories of personal data processed;

¹⁵ Paragraph 10.4 of the Proposed Guidelines.

¹⁶ Paragraphs 7.5 and 8.3 of the Proposed Guidelines.

¹⁷ Question 4 of the Public Consultation Questions in the Proposed Guidelines.

- (ii) retention policies;
 - (iii) data residency arrangements;
 - (iv) security standards and access controls; and
 - (v) metrics relating to the likelihood of data leakage.
- (c) A comparable approach is reflected in the EU Code of Practice for General-Purpose AI Models, under which providers must supply downstream integrators with the documentation they require within a defined period.¹⁸
- (d) Establishing a regulatory baseline of this nature would meaningfully reduce the information asymmetry that System Deployers presently face, especially those with limited bargaining power relative to large foundation model providers.
- (e) Consistent with the foregoing, ASIFMA recommends that the Proposed Guidelines expressly recognise that:
- (i) where a System Deployer has made reasonable, risk-based inquiries into the System Provider’s data protection measures before entering into the contract;
 - (ii) the upstream System Provider does not provide or provides incomplete information or documentation regarding its data protection measures to the System Deployer; and
 - (iii) in response, the System Deployer has made reasonable efforts to implement compensating controls on its end to protect the personal data notwithstanding its lack of visibility on the upstream System Provider’s data protection measures; then
 - (iv) the System Deployer should be given due credit and not be treated as being non-compliant with its PDPA obligations.
18. On the topic of data protection safeguards, ASIFMA would also like to take this opportunity to seek clarification on the following paragraphs of the Proposed Guidelines.
- (a) In paragraph 5.1, organisations are encouraged to practice data minimisation when developing Generative AI Models to minimise unnecessary risks.
- We would appreciate greater clarity on what the Commission means by “data minimisation” in this context. This is to avoid any misunderstanding of the concept as requiring that only minimal volumes of data be collected, since this would undermine the performance and reliability of Generative AI Models.
- (b) In paragraph 9.4, System Deployers are encouraged to develop clear written policies and document processes in relation to the safeguards undertaken, with the Proposed Guidelines noting that an organisation may demonstrate accountability in compliance with the PDPA by pre-emptively making such policies available (e.g. on its website).
- We would appreciate confirmation that System Deployers are not expected to publish policies on their websites demonstrating how they intend to safeguard personal data in the deployment of their AI systems if the AI Systems deployed are purely internal tools for operational efficiency which do not make decisions affecting individuals.

¹⁸ Measure 1.2 of EU Code of Practice for General-Purpose AI Models – Transparency Chapter.

19. Finally, ASIFMA recommends that the Proposed Guidelines be amended to make clear that the roles of Model Provider, System Provider and System Deployer are not fixed and/or mutually exclusive designations. Instead, it is possible that a single entity may potentially take on more than one of such functional roles at the same time.
- (a) In practice, a single entity frequently moves between roles, depending on the function it performs at a given stage of the Generative AI lifecycle. For instance, a System Deployer that deploys a Generative AI System in-house may take on the function of a System Provider if it also designed and developed that AI System.¹⁹
 - (b) ASIFMA submits that the Proposed Guidelines' present structure (i.e. framing responsibilities according to the three categories of Model Provider, System Provider and System Deployer) has created some uncertainties over whether the three roles are treated as mutually exclusive categories into which every AI entity must be placed.
 - (c) Certain paragraphs in the Proposed Guidelines appear to perpetuate this uncertainty and may even appear contradictory with each other when read. For example:
 - (i) Paragraph 8.1 states that the section on System Providers "excludes" entities that develop and "deploy" Generative AI Systems in-house.
 - (ii) Yet, paragraph 9.1 on System Deployers states that entities which "develop" and deploy AI in-house may take on the roles of Model and/or System Providers.
 - (d) We therefore recommend that the Proposed Guidelines make clear that:
 - (i) the three roles should be characterised as functions that are not mutually exclusive;
 - (ii) the obligations and safeguards expected of an AI entity follow the function or functions it performs in a given data processing activity; and
 - (iii) an entity performing more than one function is expected to comply with the responsibilities attaching to each function it performs.
 - (e) It would further assist for the Proposed Guidelines to provide additional examples illustrating the distinction between a Model Provider, a System Provider and a System Deployer, and the distinction between a Generative AI Model and a Generative AI System.
 - (f) Relatedly, ASIFMA invites the Commission to consider whether an entity that merely makes a Generative AI Model or System available, but plays no part in its development, should properly be treated as a "provider" at all, since this appears to be how the term "Model Provider" and "System Provider" have been defined.²⁰
 - (i) Such an entity is unlikely to be able to control the technical characteristics of the model or system, including how personal data is processed, and may therefore be unable to discharge the data protection responsibilities that the Proposed Guidelines contemplate for providers.

¹⁹ Paragraph 9.1 of the Proposed Guidelines.

²⁰ Paragraph 6.2 of the Proposed Guidelines.

- (ii) Consistent with the approach under the EU AI Act, which treats such entities as “distributors”,²¹ the Commission may wish to consider whether the definitions at paragraph 6.2 should distinguish entities that develop a model or system from those that merely make it available.

Question 5: Are there other agent-specific data challenges or risks that would be helpful to clarify at paragraph 9.5?

20. ASIFMA submits that the following data protection challenges and risks are worth clarifying in the Proposed Guidelines:
- (a) entities (including AI agents) accessing and/or using personal data beyond their defined purpose;
 - (b) entities (including AI agents) retrieving more personal data than is necessary for a given task;
 - (c) personal data being passed to third-party tools (including by or to AI agents) outside the deployer's control;
 - (d) entities (including AI agents) retaining data from prior sessions for longer than necessary; and
 - (e) personal data flowing across multiple entities (including AI agents) without clear accountability.
21. While ASIFMA agrees that the data protection risks associated with agentic functionalities merit attention, we propose that those risks would be most durably addressed if framed in a technology-neutral, outcome-focused manner (as above), rather than by reference to whether an AI system is characterised as “agentic”.
- (a) This is because, as the Proposed Guidelines acknowledge at paragraph 9.5, the definition of AI agents remains unsettled.
 - (b) Framing the guidance in a technology-neutral manner would make it more durable as such AI systems evolve, and also easier for organisations to operationalise.
22. Finally, ASIFMA notes that paragraph 9.5 of the Proposed Guidelines requires System Deployers to regularly review the sufficiency of their safeguards, especially if their Generative AI Systems have agentic functionalities. It would assist if the Commission can confirm whether a periodic, risk-based programme of control evaluation and risk assessment would be regarded as sufficient, or whether continuous or real-time monitoring is expected.

C. Conclusion

23. ASIFMA welcomes the Commission’s initiative in developing the Proposed Guidelines and supports the objective of enabling the responsible use of personal data in Generative AI while safeguarding the interests of individuals. Our members are committed to working constructively with the Commission to that end.

²¹ Article 3(7) of the EU AI Act.

24. ASIFMA is grateful for the opportunity to contribute to this consultation and would welcome the opportunity to discuss any aspect of this response with the Commission. Should the Commission require any further information, or if ASIFMA may be of assistance during this process, please do not hesitate to contact **Rishi Kapoor** through the channels stated on the cover page of this document.