PERSONAL DATA
PROTECTION COMMISSION
SINGAPORE

RESPONSE TO FEEDBACK ON PUBLIC CONSULTATION ON PROPOSED ADVISORY GUIDELINES ON USE OF PERSONAL DATA IN GENERATIVE AI

Issued 20 July 2026

Supported by:

INFOCOMM MEDIA DEVELOPMENT AUTHORITY

TABLE OF CONTENTS

PART I: INTRODUCTION

1 Background

1.1 On 2 June 2026, the Personal Data Protection Commission ("**Commission**") launched a public consultation on its proposed Advisory Guidelines on Use of Personal Data in Generative AI ("**Guidelines**").

1.2 During the consultation, the Commission sought views on the Guidelines for situations where the development and deployment of Generative AI involve the use of personal data in scenarios governed by the Personal Data Protection Act 2012 ("**PDPA**"). These include:

a) The collection and use of personal data to develop Generative AI;

b) The allocation of personal data protection responsibilities across the Generative AI lifecycle; and

c) The handling of individuals' requests concerning the processing of their personal data for Generative AI.

1.3 The public consultation closed on 1 July 2026 with 40 responses from organisations representing various sectors (e.g. technology, finance, healthcare, entertainment) and three individuals. Please refer to the Commission's website for the full list of respondents and their submissions.[1] The Commission thanks all respondents for the comments submitted.

1.4 This note summarises the key matters raised in the public consultation and provides the Commission's responses.

PART II: OVERVIEW OF ISSUES

2 Publicly Available Exception

2.1 Respondents welcomed clarification on how the Publicly Available Exception applies to personal data collection for Generative AI Model development. In addition to the examples cited, respondents sought the Commission's views on whether measures including token-based access, API rate limits, age gates, geo-blocking and robots.txt are digital barriers. They also requested examples where personal data behind digital barriers would likely not be publicly available.

[1] See https://www.pdpc.gov.sg/organisations/regulations-decisions/public-consultations/public-consultation-on-the-proposed-advisory-guidelines-on-use-of-personal-data-in-generative-ai

2.2 On the Commission's recommended best practice that collecting organisations notify controlling organisations of their intention to web-scrape, some respondents noted that this could delay Generative AI development and would be practically challenging to implement (e.g. assessing who in the controlling organisation should be notified, or whether organisations should wait a reasonable period before proceeding). They cautioned that controlling organisations could use the opportunity of being notified to impose additional restrictions on data access, which would undermine reliance on the Publicly Available Exception.

2.3 The Commission has enhanced its guidance for organisations by expanding its list of examples of digital barriers and including illustrative box stories. Having considered respondents' feedback, the Commission has removed the best practice of notifying controlling organisations from the Guidelines. Instead, the Commission has included the following clarifications.

2.4 First, where it is reasonably arguable that data behind a digital barrier is not publicly available, but an organisation nevertheless relies on the Publicly Available Exception, it must explain its assessment and reasoning in a Data Protection Impact Assessment or other written record(s). It must also produce this documentation if the Commission requires. This will ensure that organisations properly assess whether data is publicly available and keep proper records. Second, where organisations and/or individuals make personal data available online, they should implement appropriate digital barriers where they do not intend, or have not obtained consent, for their data to be web-scraped. This will encourage the use of clear safeguards to restrict access to personal data online.

3 Consent and Notification Obligations

3.1 Most respondents were supportive of the intent of AI-Specific Notifications. Respondents sought confirmation that the Consent and Notification Obligations only apply in cases where deemed consent and exceptions to the Consent Obligation are not applicable. There were also calls for explicit guidance on how the business improvement and research exceptions to consent may apply to Generative AI.

3.2 Respondents explained that organisations already use various mechanisms to notify individuals of and obtain consent for the use of User Data, including privacy policies, terms of service, in-product notifications and centralised resource hubs. They shared that requiring organisations to provide the information listed at paragraph 4.6 of the Guidelines through a standalone "AI-Specific section" or notification would create an unnecessary compliance burden, particularly where such information was already available through existing mechanisms.

3.3 To the Commission's query on whether AI-Specific Notifications should be provided

in cases beyond Generative AI Model training and/or fine-tuning, respondents were divided. Some responses preferred AI-Specific Notifications for a narrower range of circumstances (i.e. specifically where User Data is used for general-purpose model development). Others felt that AI-Specific Notifications may be helpful in broader circumstances (e.g. where the use of User Data materially affects individuals) but also cautioned against notifications for "routine" processing (e.g. fraud detection, records summarisation).

3.4 The Commission has made it clearer that the Consent and Notification Obligations only apply where organisations determine that deemed consent and exceptions to the Consent Obligation are not applicable. Such exceptions include the business improvement and research exceptions. In its Advisory Guidelines on Use of Personal Data in AI Recommendation and Decision Systems, the Commission has already given guidance on the applicability of these exceptions, which continues to be relevant for Generative AI.

3.5 The Commission has also clarified that "AI-Specific Notifications" is a shorthand for the provision of sufficient information on the use of User Data for Generative AI, to enable meaningful consent. It does not require organisations to abide by a list-style template or standard notification format. Provided that the information provided enables individuals to understand what User Data is being used and how, organisations retain discretion to determine the appropriate manner and form of notification, including supplementing or directing individuals to their existing privacy policies, terms of use etc.

3.6 The Commission has weighed the various positions on when AI-Specific Notifications are appropriate. While the Commission's view is that these are most critical in cases of training and/or fine-tuning Generative AI Models, it is supportive of organisations that choose to provide AI-Specific Notifications in respect of their downstream use cases, to enhance transparency for end-users. It has also clarified that "training" and "fine-tuning" refer only to activities that develop or modify a model's parameters or underlying capabilities. It does not include inference or operational activities such as Retrieval-Augmented Generation.

4 Data Protection Responsibilities of Generative AI Stakeholders

4.1 Respondents were supportive of the Commission's mapping of PDPA obligations to key Generative AI stakeholders (i.e. Model Providers, System Providers and System Deployers). However, several noted that an entity's role may vary depending on the context of data processing, and entities may perform multiple roles simultaneously. These respondents sought clearer guidance on how to determine which role(s) apply to them and, where they perform multiple roles, how to reconcile any overlapping

obligations.

4.2 On what additional information would be helpful for Model and System Providers to share with their downstream stakeholders, key suggestions were to include incident response and data breach notification procedures, sub-processor and/or vendor lists and cross border transfer safeguards. Some respondents also called for the development of a standardised disclosure format. Model and System Provider respondents disagreed with overly granular disclosures that would expose sensitive technical safeguards or be operationally burdensome.

4.3 Respondents provided insightful comments on agent-specific data challenges. They explained the heightened risks of agentic deployments to include persistent memory as a distinct source of personal data, unauthorised data disclosure through external connectors and tools and unclear responsibility allocation for multi-agent systems. Respondents encouraged the Commission to provide guidance on how organisations should mitigate these risks.

4.4 The Commission recognises that Generative AI stakeholders can hold multiple roles and has explained that, in such cases, stakeholders must have policies and practices as necessary to meet their various obligations. It has also added to the categories of information that Model and System Providers are encouraged to share with their downstream stakeholders.

4.5 The Commission has cited additional agent capabilities and related risks, to recognise the breadth of respondents' comments on this matter. That said, the Commission's view is that a comprehensive treatment of agent-specific data issues, which are still emerging and complex, is beyond the scope of the Guidelines. These issues will be the subject of separate study and guidance.

5 Access and Correction Obligations

5.1 Most respondents agreed with the Commission's proposed best practices to support compliance with the Access and Correction Obligations, while observing that their applicability would depend on specific contexts. For example, it would be impractical to expect Model Providers, who process massive datasets from highly heterogenous sources, to document a comprehensive lineage of training data. Given the challenges of removing inaccurate data from trained models and systems, respondents also recommended the use of output safeguards. A few responses noted that machine unlearning is still an emerging field and it would be premature to cite it as a measure for organisations to track and adopt.

5.2 There were questions on whether model parameters and inference data may also be subject to access and correction requests. Some respondents noted that, depending

on the type(s) of data concerned, Generative AI stakeholders may be constrained in responding to individual requests. System Providers, for instance, would be unable to provide access to model training data, which they typically do not have possession of or control over. Guidance was sought on how stakeholders should approach their responsibilities in these circumstances.

5.3 The Commission has made relevant refinements to its list of best practices. It has also clarified that the Access and Correction Obligations apply to personal data which has been collected, used or disclosed for model or system development or deployment, and outlined considerations for different stakeholders in responding to individual requests.

PART III: CONCLUSION

6.1 The Commission will continue to assess the need to provide guidance on the use of personal data in additional AI contexts through advisory guidelines, technical guides or other resources to assist organisations in meeting their obligations under the PDPA. Organisations should visit www.pdpc.gov.sg for more information.

6.2 Once again, the Commission thanks all respondents for their comments and participation in the public consultation.

END OF DOCUMENT