20 July 2026

Personal Data Protection Commission of Singapore
10 Pasir Panjang Road, #03-01
Mapletree Business City, Singapore 117438
Submitted via the PDPC public consultation portal

Re: Comments of the Association for Competitive Technology (ACT) on the Proposed Advisory Guidelines on Use of Personal Data in Generative AI (issued 2 June 2026)

Dear Sir or Madam,

I. Introduction and interest of ACT

The Association for Competitive Technology (ACT)[1] welcomes the opportunity to comment on the Personal Data Protection Commission's (the Commission's) Proposed Advisory Guidelines on Use of Personal Data in Generative AI (the Proposed Guidelines). ACT represents the global small business technology developer community. Our members build the applications and connected products that increasingly rely on generative artificial intelligence (AI), and they do so as small and medium-sized enterprises (SMEs) operating across the development, deployment, and post-deployment stages that the Proposed Guidelines address.

Singapore is an important market and a global reference point for pragmatic, innovation-friendly data governance. Many of our members deploy generative AI systems into Singapore or build on models and platforms that serve Singapore users. The Commission's decision to clarify how the existing Personal Data Protection Act 2012 (PDPA) applies to generative AI, rather than to erect a new and separate regulatory regime, reflects an approach that ACT has consistently urged upon regulators worldwide. We commend the Commission for grounding this guidance in established obligations and for issuing it in advisory form.

We are especially encouraged that the Proposed Guidelines allocate responsibilities across a defined set of stakeholders, namely Model Providers, System Providers, and System Deployers. This structure closely mirrors ACT's own Roles and Interdependencies Framework, a consensus taxonomy developed with our small business community that maps the actors across the AI value chain and assigns responsibility to those best placed to manage a given risk. We offer the comments below in that shared-responsibility spirit, with the aim of making the final guidance clear, proportionate, and workable for the SMEs that make up much of Singapore's and the world's AI economy.

II. General observations

[1] Contact: Brian Scarpelli, Senior Global Policy Director, bscarpelli@actonline.org

Before turning to the Commission's five specific questions, we offer four cross-cutting observations that inform our detailed responses.

The observations that follow draw on ACT's published work on AI governance, developed with and for our small business developer community and offered here as background. These include ACT's AI Policy Principles, which set out our risk-based, interoperable approach to AI regulation, and ACT's AI Roles and Interdependencies Framework, a consensus taxonomy that maps responsibilities across the AI value chain and directly parallels the Model Provider, System Provider, and System Deployer structure of the Proposed Guidelines. A fuller record of ACT's global regulatory engagement is available in our publications library.

1. Proportionality and calibration to demonstrated risk.

ACT's foundational AI policy principle is that obligations should scale with the size of the actor and the risk of the application, and that the level of review, assurance, and oversight should be proportionate to demonstrated harms. A small developer fine-tuning a model on a modest dataset does not present the same risk profile as a large frontier developer training on internet-scale data. Where the Proposed Guidelines articulate best practices, we ask the Commission to make clear that these are scalable expectations, to be met through means reasonable for the actor's size and role, and not de facto obligations that fall hardest on the smallest participants.

2. Interoperability with global frameworks.

Our members operate globally and cannot maintain a separate compliance architecture for each jurisdiction. The Proposed Guidelines already draw on concepts familiar from other regimes, including the value-chain allocation of roles found in the EU AI Act and the risk-management vocabulary of the U.S. National Institute of Standards and Technology AI Risk Management Framework. We encourage the Commission to keep notification content, transparency requirements, and stakeholder definitions interoperable with these frameworks and with the accountability principle familiar from the GDPR, so that an SME can satisfy Singapore's expectations using substantially the same notices, records, and controls it maintains elsewhere.

3. Preserving lawful access to data for AI development.

Access to diverse, high-quality data is the foundation of competitive AI development, and it is precisely where smaller developers are most disadvantaged relative to large incumbents. ACT supports robust privacy protection, including respect for the access controls and preferences that individuals and organisations express through technical means, and we

also caution against guidance that would, in practice, restrict lawful data processing more tightly than the PDPA requires. The Commission's recognition of the Publicly Available Exception, including for certain data behind digital barriers, is a welcome and balanced starting point that we address further below.

4. Shared responsibility follows control and capacity.

Responsibility for mitigating a given risk should rest with the actor that has the knowledge and the practical ability to address it, together with appropriate incentives to do so. This principle should guide how the Commission allocates duties among Model Providers, System Providers, and System Deployers, and it argues strongly for upstream actors to furnish downstream SMEs with the information those SMEs need to meet their own obligations. We return to this point in our response to Question 4.

III. Responses to the consultation questions

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

ACT supports the Commission's balanced treatment of the Publicly Available Exception and its recognition that the mere presence of a digital barrier does not, by itself, remove data from the scope of the exception. Additional examples would aid SMEs, which rarely have the legal resources to resolve close questions on their own. We suggest the Commission address the following categories:

- **Machine-readable access-preference signals.** The Robots Exclusion Protocol (robots.txt) and emerging standards that let a site express AI-training preferences in a machine-readable form are becoming central to responsible data collection. Clarifying how such signals bear on the digital-barrier analysis, and confirming that honouring a published machine-readable signal is a reasonable and sufficient step, would give small developers a clear and low-cost path to compliance.
- **Technical throttling and volume limits.** Rate limiting, IP-based access limits, and similar measures restrict how data is accessed without necessarily restricting who may access it. Guidance on whether, and when, these constitute meaningful restrictions would be valuable.
- **The boundary between contractual terms and digital barriers.** We strongly welcome paragraph 3.10, which distinguishes the availability of the exception under the PDPA from separate contractual and general-law obligations. Reinforcing that a term of service prohibition, standing alone, is a contractual matter and not itself a digital barrier

for PDPA purposes would prevent SMEs from over-reading terms of use as data-protection prohibitions. ACT encourages organisations to respect applicable contractual terms and technical access controls as a matter of responsible practice, and guidance on the interaction between terms of use and PDPA compliance would be welcome..

We also offer one caution on the best-practice notification set out at paragraphs 3.7 to 3.9. Requiring, even as best practice, that a collecting organisation notify each source organisation before scraping publicly available data behind a barrier could become disproportionately burdensome for a small developer that draws on many sources. We ask the Commission to confirm that this best practice is scalable and can be satisfied through reasonable, standardised means, such as honouring published machine-readable access signals, rather than through individualised outreach to every source. The Commission's dispute-resolution role at paragraph 3.9 is helpful, and we suggest the final guidance make clear that a collecting organisation acting in good faith on a reasonable public-availability assessment should not be penalised while such a dispute is pending.

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

ACT supports meaningful transparency and informed consent, and our AI policy principles call expressly for explicit communication of allowable use together with periodic review of informed consent. We agree with the Commission that general notifications citing only "new product development" are an insufficient basis for large-scale training on User Data, and that clear AI-specific notification is appropriate at the training and fine-tuning stage where the material and least-anticipated use of personal data typically occurs.

We would caution, however, against extending mandatory AI-specific notification broadly to inference and ordinary deployment. Requiring a distinct AI-specific notice at every interaction risks notification fatigue that dulls rather than sharpens individual understanding, and it imposes cost on small deployers without a commensurate privacy benefit where existing layered privacy notices and the Purpose Limitation Obligation already govern the processing. We recommend the following calibrated approach:

- Retain AI-specific notification as the standard for training and fine-tuning, the stage at which it does the most work.
- For other stages, rely on layered general privacy notices combined with purpose limitation, and reserve additional targeted notification for cases involving personal data of a more sensitive nature or genuinely unanticipated secondary uses, which is

consistent with the risk-based calibration the Proposed Guidelines already reflect at paragraph 4.5.

- Keep the required content of AI-specific notifications aligned with the information categories used under the GDPR and comparable regimes, so that a single notice architecture can serve an SME across markets.

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)?

On the substance of notifications (paragraph 4.6). We endorse the Commission's practical, plain-language orientation and offer three additions:

- **Layered notices.** Confirm that a short, prominent summary linked to a fuller explanation satisfies the guidance. Layering serves comprehension and is well suited to small-screen and in-product contexts.
- **Notification by reference.** Permit a System Deployer to satisfy its notification role by pointing clearly to accurate upstream documentation from a Model or System Provider, rather than re-authoring technical descriptions it is not best placed to write. This reflects the shared-responsibility principle and reduces duplicative effort for the smallest actors.
- **Standardised, machine-readable formats.** Encourage, without mandating, common templates so that notices are comparable across products and cheaper for SMEs to produce and for individuals to understand.

On access and correction (paragraph 10.4). We appreciate the Commission's candid recognition of the technical realities of embeddings, temporary context windows, and the current immaturity of machine unlearning. We suggest the final guidance make two points explicit:

- **Output-level and retrieval-level remediation can be reasonable.** Where retraining or unlearning is technically infeasible or disproportionately costly, correcting or suppressing personal data at the output layer, or in a Retrieval-Augmented Generation database, should be recognised as a reasonable and proportionate means of responding to a request. Tying the standard to reasonableness, rather than to a specific technical method, keeps the guidance durable as techniques mature.
- **Privacy Enhancing Technologies as a pathway.** Cross-reference the Commission's own Privacy Enhancing Technologies Adoption Guide and Guide to Basic

Anonymisation as concrete, proportionate tools that help SMEs both minimise personal data upstream and respond to requests downstream.

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?

This question sits at the heart of ACT's Roles and Interdependencies Framework, and we regard it as the single most important opportunity in the Proposed Guidelines to help SMEs. A System Deployer, which the Proposed Guidelines rightly place under primary responsibility for PDPA compliance, can only discharge that responsibility if the Model and System Providers upstream give it the information it needs. Small deployers are least able to reverse-engineer upstream practices, so clear, standardised upstream disclosure is both a fairness measure and a practical enabler of compliance.

We recommend the Commission encourage Model and System Providers to provide a concise data-protection transparency policy document that may usefully cover:

- Whether and what categories of personal data were used in training or fine-tuning, at a level of generality that is informative without compromising security, intellectual property, or proprietary information.
- Data residency, retention, and deletion practices for data processed during inference.
- Access controls, input and output filtering, and any Privacy Enhancing Technologies applied.
- Testing and performance information relevant to data protection, such as measured rates of personal-data leakage or redaction.
- Known limitations and residual risks that a downstream deployer should account for.

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

ACT welcomes the Commission's attention to agentic systems and its sensible reference to the IMDA Model AI Governance Framework for Agentic AI, which supports interoperability and avoids a divergent Singapore-specific track. Because the definition of an AI agent remains unsettled, we encourage the Commission to keep this section principles-based. Prescriptive rules written to today's agent architectures are likely to date quickly, a concern ACT has raised consistently as AI technology evolves. Within that principles-based posture, the following agent-specific issues would benefit from clarification:

- **Accountability across tool and agent hand-offs.** When an agent invokes third-party tools or other agents, responsibility for the resulting personal-data processing should follow control and the capacity to mitigate, consistent with the shared-responsibility principle rather than defaulting to the smallest or most visible actor in the chain.
- **Persistent agent memory.** Long-lived agent memory raises retention and data-minimisation questions distinct from a single-turn context window. Guidance on treating persistent memory as a discrete retention concern, subject to review and purposeful limitation, would help deployers scope it responsibly.
- **Blurring of control and user data planes.** As the Proposed Guidelines note in the System Provider context, agentic operation can obscure the boundary between instructions and data, which heightens prompt-injection and data-exfiltration risk. Cross-referencing that discussion at paragraph 9.5 would give deployers a coherent view.
- **Human oversight and escalation as a proportionate safeguard.** We support the Commission's recognition of escalation protocols for complex or high-risk tasks. Framing human-in-the-loop review as a proportionate, scalable safeguard, rather than a uniform requirement, lets SMEs match oversight to the risk of the specific use case.

IV. Conclusion

ACT commends the Commission for guidance that clarifies existing obligations, allocates responsibility across a sensible value-chain taxonomy, and does so in advisory form. With the calibrations described above, above all a consistent emphasis on proportionality, interoperability, and upstream transparency that equips downstream SMEs to comply, the final guidance can protect individuals effectively while keeping generative AI development open to the small businesses that drive competition and innovation. We would welcome the opportunity to serve as a resource to the Commission as this work proceeds, and we thank you for your consideration of these views.

Respectfully submitted,

Brian Scarpelli
Senior Global Policy Counsel
Association for Competitive Technology (ACT)
Washington, DC