AI and personal data: a deployer's checklist

For firms that use generative AI tools and agents in day-to-day work. Based on the Singapore Government's Personal Data Protection Commission (PDPC) Advisory Guidelines on Use of Personal Data in Generative AI (20 July 2026) and its Advisory Guidelines on Use of Personal Data in AI Recommendation and Decision Systems (1 March 2024). The guidelines explain how the Personal Data Protection Act applies; they do not add to it. The firm that switches a tool on carries primary responsibility for how it handles personal data, whoever built it.

1. Before you buy or switch on a tool

- ☐ **Training.** Confirm in writing whether the vendor uses your inputs, documents or outputs to train or fine-tune a model. Choose the no-training setting where one exists.
- ☐ **Residency.** Know where prompts, uploads and outputs are stored and processed, and whether a Singapore or regional option is available.
- ☐ **Retention.** Know how long the vendor keeps prompts, outputs and uploaded files, and how to shorten it or delete them.
- ☐ **Access.** Know who at the vendor can see your data, and what access controls and encryption apply.
- ☐ **Breach procedure.** Know how and how fast the vendor will tell you about a security incident. The PDPA requires you to assess a suspected breach within 30 days and to notify the PDPC within three calendar days of finding it notifiable.
- ☐ **Published safeguards.** Keep a copy of the vendor's data protection documentation (security sheet, data processing terms, system or model card). This is the "sufficient information on upstream safeguards" the PDPC expects you to hold.

2. Setting the tool up

- ☐ **Purpose.** Write one sentence stating what the system is for. Everything else follows from it.
- ☐ **Data needed.** List the personal data the purpose genuinely requires. Leave out the rest.
- ☐ **Access limits.** Configure the system so it cannot read data outside its purpose: no connection to HR, payroll, billing or customer contact lists unless the job needs them.
- ☐ **Owner.** Name the unit or person responsible for the system's prompts, outputs and logs.
- ☐ **Sensitive data.** Identify any children's, health, financial or NRIC data in scope and decide whether it should enter the system at all. If it must, apply stricter handling.

3. Day-to-day use by staff

- ☐ **Staff rule.** Give staff a short written rule on what personal data may and may not go into a prompt or upload, with examples from your own work.
- ☐ **Account hygiene.** Use the firm's business accounts, not personal logins, so the firm's settings and retention apply.
- ☐ **Identifier scanning.** Where the tool allows, scan or redact common identifiers (names, NRIC or FIN numbers, phone numbers, account numbers) before they reach the model.

- ☐ **Log access.** Limit who can read prompt and output logs to named people who maintain the system.
- ☐ **Output checks.** Decide which outputs a person must review before they are acted on or sent to a customer.

4. Agents that act, not just answer

- ☐ **Scope.** Define exactly which systems, files and tools the agent may reach. Grant access by role, not by default.
- ☐ **Classification.** Tag sensitive personal data so the agent handles it under stricter rules or not at all.
- ☐ **Escalation.** Set the conditions under which the agent must stop and hand the task to a person: large amounts, unusual counterparties, anything outside its brief.
- ☐ **Activity log.** Keep a record of what the agent read, decided and did, and who can inspect it.
- ☐ **Framework.** For higher-risk agent use, read IMDA's Model AI Governance Framework for Agentic AI (January 2026).

5. Policies and records

- ☐ **AI governance statement.** Write a short internal document covering purpose, data, owners, staff rule, vendor safeguards and review cycle. Circulate it to all staff.
- ☐ **Published version.** Consider a plain-language summary on your website. The PDPC says this demonstrates accountability, and it answers most vendor assessment questions in advance.
- ☐ **Privacy notice.** Check that your privacy notice covers the AI uses you have introduced. If customer data will be used to train a model, a general "product improvement" line is not sufficient: name the data, the use and how to opt out.
- ☐ **Data Protection Officer.** Confirm the DPO is registered with the PDPC and knows which AI systems are in use.
- ☐ **Access and correction.** Make sure a customer's request to see or correct their data reaches any retrieval database or uploaded files the AI system uses, not only the main records.
- ☐ **Review date.** Set a date to review the safeguards, and bring it forward whenever the tool gains new capabilities.

Quick reference: the four questions to ask any AI vendor

1. Do you train on our data, and can we turn it off?
2. Where is our data stored, and for how long?
3. Who at your end can see it?
4. How and when do you tell us about a breach?
