



QUAY INSIGHTS

February 2025

The Privacy Commissioner's guidance on generative AI: Part I – Guidance for AI Developers

The Privacy Commissioner has published two sets of guidelines relevant to generative AI models: *Guidance on privacy and developing and training generative AI models (Guidance for AI Developers)*, targeted at developers of AI products, and *Guidance on privacy and the use of commercially available AI products (Guidance for AI Deployers)*, targeted at the deployers and users of AI. Each guide details how the Australian Privacy Principles (APPs) in the Privacy Act apply to the contexts that arise when developing AI models and systems and when using AI.

This Quay Insight summarises the Privacy Commissioner's Top 5 Takeaways in the Guidance for AI Developers.

Top 5 privacy takeaways for AI Developers

The Privacy Commissioner presents 5 top takeaways to AI developers in meeting their privacy obligations:

1. **Accuracy** – AI developers must take reasonable steps to ensure accuracy in generative AI models, commensurate with the likely increased level of risk in an AI context, including by using high quality datasets and undertaking appropriate testing.

This reflects the APP 10 requirement to take reasonable steps to ensure that the personal information the entity collects, uses and discloses is accurate, up-to-date and complete, and (for use and disclosure) is relevant. The guidance highlights that one such reasonable step may be to place disclaimers on the AI product, clearly and specifically communicating any limitations in its accuracy.

2. **Legal basis for collection and use** – The guidance emphasises that “just because data is publicly available or otherwise accessible does not mean it can be legally used to train or fine tune generative AI models or systems.” To ensure compliance AI developers must correctly identify whether the information they collect is personal information under the Privacy Act, and may need to consider additional steps such as deleting personal information collected.

This serves as a reminder to AI developers that, under APP 3, non-sensitive personal information must only be collected where reasonably necessary for one or more of the entity's

This publication is not intended to be comprehensive on the topics with which it deals. It is not intended to be relied upon or provide legal advice on the topic. Specific professional advice should be sought.

www.quaylaw.com

Level 32, 180 George Street, Sydney NSW 2000, Australia

LIABILITY LIMITED BY A SCHEME APPROVED UNDER PROFESSIONAL STANDARDS LEGISLATION

functions or activities, and only by lawful and fair means. It also highlights collection from the internet, and other publicly available sources, is a “collection” of personal information necessitating consideration of privacy obligations.

3. **Sensitive information** – AI developers must take particular care with sensitive information which generally requires consent to be collected. The guidance notes that many images or videos of individuals contain sensitive information and may therefore require consent before being scraped from the internet or collected from a third-party dataset.

This reflects APP 3.3, which prohibits the non-consensual collection of sensitive information except in very limited circumstances. There are a number of categories forming part of the definition of “sensitive information” under the Privacy Act. Examples include personal information about an individual’s race, ethnicity, political opinion, religion, sexual orientation, criminal record, health, genetics and certain kinds of biometric information.

4. **Repurposing data holdings** – AI developers seeking to use personal information from their own data holdings to train an AI model, where AI was not the primary purpose of collection, will need to carefully consider their privacy obligations.

The guidance connects with the requirement in APP 6 to generally only use personal information for the primary purpose for which it was collected. APP 6 permits use for a secondary purpose where the individual would reasonably expect it and the secondary purpose is *directly related* to the primary purpose in the case of sensitive information, or *related* to the primary purpose for non-sensitive personal information.

5. **Consents and opt-out is best regulatory practice** – The guidance states that where an AI developer is considering using personal information collected for a non-AI related purpose, and cannot clearly establish that the exception referenced above applies, to avoid regulatory risk, it should seek consent from the individual and offer them a meaningful and informed ability to opt out.

What should AI Developers do?


The Guidance for AI Developers provides valuable insight into how the Privacy Commissioner intends to apply the APPs in the event that a particular AI model is investigated by the regulator. The Privacy Commissioner has spoken publicly about addressing AI harms as a priority and has stated that the OAIC is developing a regulatory strategy in the field of generative AI. Given these comments and the recent publication of the guidance, AI developers are on notice of the potential for regulatory enforcement action in this space in 2025.

Contacts




Cate Cloudsdale
Counsel

Quay Law Partners
Level 32, 180 George Street,
Sydney NSW 2000
T +61 461 477 550
E cate@quaylaw.com
www.quaylaw.com



Angela Flannery
Partner

Quay Law Partners
Level 32, 180 George Street,
Sydney NSW 2000
T +61 419 489 093
E angela@quaylaw.com
www.quaylaw.com



Dave Poddar

Partner

Quay Law Partners
Level 32, 180 George Street,
Sydney NSW 2000
T +61 422 800 415
E dave@quaylaw.com
www.quaylaw.com