Someone Pasted Patient Data Into ChatGPT — What Now?

Sooner or later it happens. A calm, prepared first response — not panic, and not pretending it didn't happen — is what protects the practice and the patient.

Someone Pasted Patient Data Into ChatGPT — What Now?

> The worst time to work out what to do about an AI privacy incident is during one. Decide the steps now, while it's calm.

This is general educational material for dental practice owners and staff, not legal advice. It is a first-response guide, not a substitute for qualified privacy or legal advice when an incident occurs.

> **Two privacy laws apply in NSW.** As well as the Commonwealth *Privacy Act 1988* and its Australian Privacy Principles (APPs), dental practices in NSW are also bound by the NSW *Health Records and Information Privacy Act 2002* (HRIP Act) and its Health Privacy Principles (HPPs). Read the considerations here against both. General information, not legal advice.

Why this matters

If staff use computers, this will happen eventually — a name pasted into ChatGPT to draft a letter, a patient list dropped into a tool to "tidy it up", an AI feature found to be storing data offshore. The damage is usually made worse by two reactions: **panic**, or **pretending it didn't happen**. A short, prepared process avoids both.

This connects to the extraction cycle: once data has left the system, the question becomes what to do about the copy that's now out there.

First response — the steps

Work through these calmly. Identifying an incident is **not** the same as declaring a breach — it is gathering the facts so the right people can decide.

1. **Identify what was involved.** What information (names, health details, images, contact details)? Which tool? Which patient(s)? Roughly when?
2. **Contain it.** Stop the workflow. Where possible, delete the content from the tool, clear its history, and stop any sync. Change a password or revoke access if an account was involved.

3. **Document it factually.** Write down what happened, when, who was involved and what data — a plain record, not blame. This record matters later.
4. **Escalate.** Tell the practice owner / privacy officer straight away. One person should own the response.
5. **Assess whether it may be notifiable — with help.** The *Notifiable Data Breaches* scheme can require notifying affected people and the regulator for serious breaches. **Do not self-assess the threshold.** Whether an incident is an "eligible data breach" is a legal assessment — get qualified privacy/legal advice promptly.
6. **Act on advice without delay.** If notification is required, the scheme has timing expectations — move on advice quickly rather than sitting on it.
7. **Close the loop.** Update the staff AI policy and training so the same thing doesn't recur. Most incidents point to a gap that's easy to fix.

Two things not to do

- **Don't quietly delete and move on** as if it didn't happen — the documentation and the assessment still matter.
- **Don't ask an AI tool whether it's a notifiable breach.** That assessment requires qualified human review, not a chatbot's opinion.

Be ready before it happens

- Decide **who** the incident owner is.
- Keep this checklist somewhere staff can find it.
- Make sure staff know they should **report** an AI slip immediately — and won't be punished for owning up. The earlier it's reported, the more containable it is.
- Pair this with the Can I Paste This Into AI? guide so fewer incidents happen in the first place.

This guide is educational material only. It is not legal advice and does not determine whether any incident is a notifiable data breach. Seek qualified privacy or legal advice for a specific incident.