

# Before You Switch On an AI Receptionist or Chatbot

An AI receptionist is the most-pitched AI product in dentistry — and often the least safe to switch on first. It collects health details and stores them somewhere you may not have checked.

---

## Before You Switch On an AI Receptionist or Chatbot

An AI receptionist is the most-pitched AI in dentistry, and frequently the least safe first project. It collects patient health details and stores them — the only question is where, and who can see them.

This is general educational material for dental practice owners and managers, not legal or clinical advice. This guide is the *pre-purchase readiness* check; the [Emergency Booking and the AI Boundary](#) guide covers the urgent-patient handling boundary in depth.

**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

AI receptionists, booking bots and website chat widgets are being sold to practices as an easy win. But unlike an internal tool, this one **talks directly to patients** and **collects health information** — a patient describing a toothache, swelling or symptoms is providing health data, often after hours, into a system you may not have vetted.

That makes it a poor *first* AI project for many practices: it sits at the most sensitive point (a patient in distress sharing symptoms) and the data flows straight out to a vendor — the [extraction problem](#), but with the patient doing the typing.

### Readiness questions — ask before you sign

#### Data and storage

- Where are conversations stored, and for how long?
- Is storage or processing **overseas**? (If so, treat it as an **APP 8 / HPP** review item.)

- Who at the vendor can access the conversations? Can the practice delete them?
- Does the vendor's contract or data-processing agreement actually address this, or just reassure you verbally?

### **The clinical boundary (the big one)**

- Does the tool ever **assess urgency or symptoms**? It must **not** make clinical assessments or tell a patient whether their symptoms are serious — that is a clinician's job.
- Is there a clear, fast **escalation to a human**, especially for anything that sounds urgent?

### **Consent and notice**

- Are patients told their conversation is collected and how it's used (a privacy notice that covers this channel)?
- Does the tool collect more than it needs (full symptom descriptions when a name and callback number would do)?

### **What good looks like**

- **Don't make it your first AI project.** Start somewhere lower-stakes; come to patient-facing AI once the basics are in place.
- **Scope it tightly.** Booking and admin — not triage, not clinical advice.
- **Check the data flow before go-live**, not after: storage location, access, deletion, overseas processing.
- **Guarantee human escalation** for urgent or symptom-heavy messages.
- **Cover it in your privacy notice** and collect the minimum.

---

*This guide is educational material only. It is not legal or clinical advice. Identifying a risky workflow indicates possible exposure, not a declared breach. Seek qualified advice for your specific circumstances.*

**Disclaimer:** Educational guidance only, not legal advice. This guide is intended for practice workflow education. Do not enter patient-identifiable information into public AI tools.