An AI Vendor Just Pitched Your Practice — What to Ask Before You Say Yes

AI scribe, receptionist, recall-bot and marketing vendors are pitching dental practices every week. The demo always looks great. These are the seven questions that decide whether your practice can actually use it safely — and who's responsible when it gets something wrong.

An AI Vendor Just Pitched Your Practice — What to Ask Before You Say Yes

> This is general educational material for dental practice owners and managers, not legal advice. The regulatory points below are things to *review and confirm* with a qualified adviser for your situation — not determinations about any specific product.

Every week, another AI tool is pitched to dental practices: a scribe that writes your clinical notes, an "AI receptionist" that answers calls and books patients, a recall bot, a marketing assistant that writes your website and social posts. The demos are slick and the time-savings are real.

But the demo only answers one question: *does it work?* The questions that actually decide whether **your** practice can use it are different — and the vendor's salesperson is rarely the person who can answer them. The core principle is the one that should sit behind every AI decision in the practice:

> **Just because an AI feature exists doesn't mean your practice can safely use it.** Whatever the tool does, *you* remain responsible for patient privacy, for the clinical record, and for anything published under the practice's name.

Here is the seven-question script to run before you sign anything.

1. Where does our patient data go — and does it leave Australia?

Ask exactly where data is stored and processed, and whether any of it — or any sub-processor they use — sits **overseas**. Sending patient information to an overseas service is a cross-border disclosure question under the Privacy Act (APP 8), and it doesn't stop being your responsibility because a vendor is in the middle. Also ask the quieter question: **is our data used to train their models?** "Your data improves the product" can mean patient information becomes part of a model you can't claw back.

A good answer: clear data-residency information, named sub-processors, and a plain "no, your data is not used to train shared models."

2. Who can see it — and can you show me?

Who at the vendor can access practice or patient data, under what controls, and **can you see an access log**? A tool that can show you who accessed what, and when, is in a different league from one that can't. This is the security-and-accountability question (APP 11), and it's also the difference between "trust us" and "here's the audit trail."

A good answer: role-based access, encryption, and an audit log you can actually inspect.

3. Is it read-only, or does it *act*?

This is the single most important safety question. A tool that **reads and suggests** — and leaves a human to approve — is far safer than one that **writes, sends, books, or posts on its own**. Autonomous action is where the real risk lives: an AI that sends a message, books a patient, or publishes a post without sign-off can cause a privacy, clinical, or advertising problem before anyone notices.

A good answer: read-only by default, with explicit human approval gates before anything is sent, written, or published.

4. What happens to our data if we leave?

Ask how you **export** everything and how you get it **deleted** if you cancel — and whether they keep a copy. Health information has destruction and retention obligations (APP 11), and "we'll keep it on our servers indefinitely" is the wrong answer. (Note the flip side: your *own* clinical records still have legally required minimum retention periods — deleting the vendor's copy is not the same as deleting the record in your dental system.)

A good answer: a clean export, a defined deletion process and timeframe, and confirmation no residual copy is retained.

5. Who's liable when it gets something wrong?

AI gets things wrong. A scribe can put a wrong figure in a clinical note (a data-quality issue under APP 10), a chatbot can mishandle a patient's information, a marketing tool can publish a testimonial or outcome claim that breaches AHPRA's advertising rules. When that happens, **the practice is usually still the responsible party** — the advertiser, the record-keeper, the registered provider — not the vendor. Get the responsibilities in writing, and never assume the vendor's contract shifts the regulatory duty off you.

A good answer: the vendor is transparent that you remain the responsible party, and helps you put review steps in place rather than promising the problem away.

6. Is it — or does it act like — a medical device?

If a tool does anything that looks like triage, diagnosis, or clinical decision-making ("our AI assesses urgency", "it screens symptoms"), that can stray into **Software as a Medical Device** territory, which the TGA regulates. A booking or admin tool generally isn't a medical device; something that makes or guides a clinical judgement might be. Worth confirming before you rely on it.

A good answer: the vendor knows the distinction and can tell you, plainly, which side of it their tool sits on.

7. Can we try it without real patient data first?

The safest way to evaluate any tool is a **reversible pilot that doesn't expose patient information** — public data, synthetic/test data, read-only, easy to switch off. If a vendor can't let you trial it safely, that itself is a signal.

A good answer: a real way to pilot on non-patient data before anything live is connected.

Green flags vs red flags

Green flags	Red flags
Read-only by default; human approval before it acts	Acts autonomously; "it just handles it for you"
Australian data residency, or clear, named overseas handling	Vague or evasive about where data goes
An access/audit log you can inspect	"Trust us — it's secure" with nothing to show
Clear export + deletion, no retained copy	Won't commit to deleting your data
Transparent that *you* remain responsible	Implies the tool makes you "compliant"
Safe pilot on non-patient data	Wants live patient data on day one

The safest first AI is usually the one you control

Notice what these questions reward: tools that are **read-only, owner-approved, auditable, and don't take patient data somewhere you can't see**. That's not an accident — it's the shape of safe AI in a dental practice.

It's also why the safest *first* AI project is rarely the patient-facing one. An AI receptionist or chatbot is the highest-stakes thing to switch on. A private, read-only assistant that reads your own data and

reports back to the owner — finding the recalls and treatment plans you're leaving on the table — sits on much safer ground, because nothing happens to a patient without a human deciding.

Before you sign — and before you switch on

- Run the relevant **free public checks** first — e.g. if it's a marketing or website tool, our website and advertising scanners review what's already public.
- For the full picture across privacy, advertising, discoverability and booking, request your free **practice Blueprint** — public information only, no patient data.
- See the foundational guide, **Where Patient Data Is Protected — and Where It Escapes**, and **Is This AI Tool a Medical Device?** for the TGA question.

When you want AI working *inside* the practice that you fully control — read-only first, owner-approved, fully auditable — that's the kind of system worth building rather than renting.

This guide is educational material only. It is not legal advice and does not assess any specific product. Confirm privacy (Privacy Act / APPs, and state laws like the NSW HRIP Act), AHPRA advertising, and TGA questions with qualified advisers for your circumstances.