

AI PROCUREMENT · INFRASTRUCTURE · 2026

Not All AI Providers Are Equal. Ask These Ten

A company that operates data centres and a company that operates an API key can write the same sentence on a website. Both statements can be entirely true.

"We have AI" is compatible with five very different companies. An operator that runs its own network and platform; a platform built on hyperscaler infrastructure; a reseller with a logo on somebody else's product; a thin application over a public model API; and a specialist model provider. All five can truthfully make the same claim, and a demonstration cannot tell them apart because the demonstration is of the model. **Your call audio passes through layers**, application, middleware, model API, cloud infrastructure, and each layer is a separate company with separate policies. **The model provider's protections attach to whoever holds that contract, not automatically to you.** Zero data retention announcements in August 2026 are genuine and welcome; they do not tell you what the intermediary you are contracting with does with your audio before or after.

■ The short version

The claim	What it is compatible with
"AI-powered transcription and summaries"	Anything from a purpose-built local pipeline to a single API call to a service on another continent.
"Enterprise-grade AI security"	Usually a restatement of the <i>model provider's</i> security posture, which is not the same as the posture of the company you are paying.
"Your data is never used to train models"	Frequently true of the model provider's API. Says nothing about what the intermediary retains, logs, caches or stores on the way through.
"Australian hosted"	Accurate about the platform, and silent about the AI feature you enabled with one click, which may be processed somewhere else entirely.

■ Also worth knowing

Layer	What it is	Whose policies govern it
1. The application	The product you log into. Handles the call, holds the recording, presents the transcript.	The company you contracted with.
2. The middleware	The glue, queues, storage, logging, retries, caches. The least visible layer and frequently where copies of your audio actually sit.	Also the company you contracted with, though they may not have documented it.
3. The model API	The service that turns audio into text, or generates a response.	The model provider, under a contract <i>they</i> hold, not you.
4. The cloud infrastructure	The compute and storage everything above runs on, in a particular country.	A hyperscaler or a data centre operator, several contracts away from you.

Layer two is where the surprises live

Most procurement attention goes to layer three, because that is where the recognisable brand names are and where the published policies are strongest. But layer two is where a system keeps **queues, retry buffers, error logs containing request payloads, caches and temporary storage**, and layer two belongs to the company you are paying, is rarely documented, and is frequently the only place your raw audio is persisted at all.

■ Questions the full guide answers

- ☐ What does it mean when a phone system provider says they have AI?
- ☐ Where does my call audio actually go when a call is transcribed by AI?
- ☐ If the AI provider says they do not train on my data, is my data safe?
- ☐ What is zero data retention and does it solve the problem?
- ☐ How should I ask about data residency for AI features?
- ☐ Is it ever fine to use a small AI provider built on a public API?
- ☐ What AI obligations stay with my business rather than my provider?

This is the summary. The guide has the detail.

Every point above is worked through in full on our site, with the Australian context, the numbers and the exceptions. If you would rather talk it through, call **1300 881 662** or visit **unidenvoice.com/get-started**.