AI SECURITY · VOICE CHANNEL · 2026

The Four Calls That Cost a Business $97,000

That is not a failure of attention. It is the design. Sophisticated fraud on the voice channel is assembled from four or five entirely unremarkable conversations, which is why reviewing a sample of calls will never find it and why.

Sophisticated voice fraud is a sequence, not a call. Four unremarkable conversations assemble the credibility for one request, and no individual call looks wrong, which is why a two per cent sampled review will never find it and why pattern analysis across the whole history will. **The cost position:** more than 84,700 cybercrime reports a year in Australia; average cost per incident **$80,850**, rising to **$97,000** for medium businesses and **$202,700** above two hundred staff; AI assessed as enabling attacks at greater scale and speed. **Where AI materially helps on voice:** cross-call pattern detection, traffic anomalies and toll fraud, verification consistency, searchable transcript evidence, complete rather than sampled review, and closing the shadow channels staff invent when calls go unanswered.

The short version

Call	If reviewed in isolation, what would a reviewer note?
1. Spelling check	Nothing. A courteous thirty-second call, handled well. Possibly a positive example.
2. Process question	Nothing. A staff member answered a reasonable question helpfully. Arguably good service.
3. Credibility call	Nothing. The caller knew internal details, so of course they were treated as known.
4. The ask	**This one is catchable**, but only after the money is gone, and only if this specific call happens to fall in the two per cent sample.

Also worth knowing

Threat	Does AI on the voice channel help?	What actually addresses it
Multi-call social engineering sequence	Materially	Cross-call pattern analysis. This is the case AI is genuinely built for.
Toll fraud / system compromise for call traffic	Materially	Real-time traffic anomaly detection with somebody empowered to block.
Inconsistent identity verification	Materially	Automated challenge sequences that do not respond to urgency or flattery.
Cannot prove what was said or authorised	Materially	Searchable, timestamped transcripts with a defined retention rule.
Policy drift under pressure	Materially	Reviewing every call rather than a sample, so breaches surface as rates.
Business comms on personal channels	Materially, indirectly	Answering calls reliably so staff stop inventing workarounds.
Synthetic / cloned voice	Partly	Not detection, a callback policy on the request type, applied without exception.

The line worth reading twice

The steepest increase. **219%**, belongs to the largest organisations, the ones with security teams, budgets and mature controls. That tells you the attack side improved faster than the defence side at the top of the market. **The inference for a smaller business is not comfort.** It is that the attempts arriving at a fifteen-person company are now built with the same tooling as those aimed at organisations employing security professionals, and a proportionate response has to account for that.

Questions the full guide answers

- [ ] How does a social engineering attack on the phone actually work?
- [ ] Why can't reviewing recorded calls catch this kind of fraud?
- [ ] Can AI detect a cloned or deepfake voice on a business call?
- [ ] What should a callback verification policy actually say?
- [ ] What are shadow channels and why are they a security problem?
- [ ] What should I do first if my budget is limited?
- [ ] Which compliance dates should I have in my calendar?

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 663 222** or visit **vocphone.com/about/contact-us**.