# What does one successful outcome cost on an AI voice agent?

Seven days of continuous production calling across three anonymised outbound tenants

Callable AI · Research Note CA-2026-02 · Published 17 August 2026

| 881 | $16,638 | $18.89 | $17.11 |
|---|---|---|---|
| qualifying successful outcomes in 7 days | combined platform spend over the same window | blended cost per successful outcome | best tenant cost per successful outcome |

## Summary

Across 11 – 17 August 2026, three outbound tenants running AI voice agents produced **881** qualifying successful outcomes against **$16,638** of platform spend — a blended **$18.89 per successful outcome**. Individual tenants ranged from **$17.11** to **$44.43**.

A seven-day window was chosen deliberately. All three tenants were fully live and calling for the entire period, so no tenant is charged for days it was not operating and every figure reflects steady-state production rather than ramp-up.

The range is the finding. A 2.6x spread in cost per outcome across three tenants on the same platform, in the same week, shows that unit economics are driven by list quality, offer and script — not by the platform fee. Contact rate alone varies from 14% to 59% of dials between these tenants.

## What counts as a success

Each tenant configures one *primary* outcome for its agent — the commercial event the campaign exists to produce. For all three tenants in this window that outcome is a qualified lead handed to the client. A call is counted once, only when the post-call extractor classifies it as that primary outcome and the result survives the platform's required-field validation gate. Outcomes flagged *invalid* by that gate are excluded.

Not counted: connected conversations, positive sentiment, callbacks, transfers that dropped, or any secondary outcome key. This is the strictest available definition and therefore a conservative denominator for cost.

# Results

### Cost per successful outcome by tenant

| Tenant | Sector | Successes (7d) | Platform cost (7d) | Cost per success |
|---|---|---|---|---|
| Tenant A | Superannuation | 818 | $14,000 | $17.11 |
| Tenant B | Residential solar | 21 | $772 | $36.78 |
| Tenant C | Financial services | 42 | $1,866 | $44.43 |
| **Combined** | — | **881** | **$16,638** | **$18.89** |

All amounts in AUD. Seven-day cost is the contracted monthly platform fee apportioned at 7/30 of a month ($60,000, $3,310 and $7,997 per month respectively). Cost per success = seven-day cost ÷ successes recorded.

# Volume behind the outcomes

| Tenant | Dials | Connected conversations | Agent talk time | Successes |
|---|---|---|---|---|
| Tenant A | 317,429 | 43,451 | 1,255 hrs | 818 |
| Tenant B | 8,068 | 1,488 | 38 hrs | 21 |
| Tenant C | 9,096 | 5,334 | 75 hrs | 42 |
| **Total** | **334,593** | **50,273** | **1,368 hrs** | **881** |

A connected conversation is a call of 25 seconds or longer, the threshold at which a two-way exchange has occurred. Talk time is billed agent-minutes on connected and unconnected attempts combined.

# Conversion efficiency

### Successes per 1,000 dials

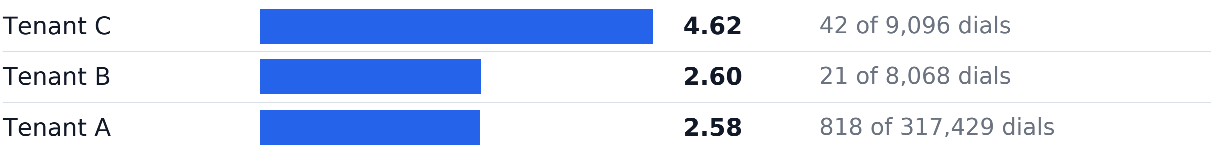


### Successes per 100 connected conversations

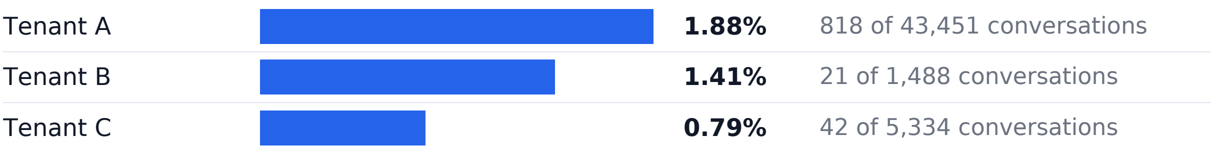


The two views disagree, and the disagreement is informative. Tenant C reaches a human on 59% of dials — by far the best contact rate in the set — and leads on outcomes per thousand dials, but converts the smallest share of the conversations it wins. Tenant A reaches a human on only 14% of dials yet converts those conversations best, and buys its absolute volume with scale: it places 18 times more dials than the other two combined. Tenant B sits between the two on contact rate and on conversion, and lands mid-table on cost per outcome.

Cost per outcome therefore tracks the product of contact rate and conversion, not contract size. The largest contract in the set is the cheapest per outcome; the middle contract is not the middle cost.

## Method

### Window

11 – 17 August 2026 inclusive, Australia/Brisbane. Seven consecutive days ending on the publication date, chosen because all three tenants were live for the whole of it. Within the week, calling activity fell on six days for Tenant A and five days for Tenants B and C; the idle days are ordinary weekend and scheduling gaps, not outages, and no day is excluded from either the numerator or the cost.

### Outcome counting

For each tenant, the agent’s configured primary outcome key was resolved from its live configuration, then post-call result records in that window matching that key were counted, excluding any record whose validation status is *invalid*. Records are one-per-call, so no call is double-counted.

### Cost basis

The contracted monthly platform fee apportioned to the window at 7/30 of a month. No tenant received a discount, credit or ramp allowance in the calculation. Unlike a calendar-month view, this basis charges each tenant only for days on which its agents were in production.

## Limitations

### Platform cost only

The denominator is platform spend. It excludes list acquisition, human follow-up, CRM licences and any commission paid on closed business. A full cost-per-acquisition figure would be higher for every tenant.

### Outcome, not revenue

A successful outcome is a qualified lead, not a closed sale. Downstream close rates differ by sector and are outside this dataset, so these figures cannot be read as cost per customer.

### One week, three tenants

A seven-day window removes ramp-up distortion but is more exposed to week-to-week variance than a monthly one, particularly for the two smaller tenants whose weekly outcome counts are in the tens. Over the preceding thirty days the same three tenants blended to $22.38 per outcome, so the weekly figure is not an outlier in aggregate, but individual tenant figures do move: Tenant B ran materially below its monthly average this week and Tenant A above it.

### Classifier dependence

Outcomes are assigned by an automated post-call extractor with a required-field validation gate, not by human review of every call. The gate is designed to fail closed — calls lacking captured qualifying data are downgraded out of the success count — so residual error biases the success count downward and the cost per outcome upward.

## Anonymisation

Tenants are identified by letter and sector only. No company name, agent name, contact record, phone number or transcript excerpt appears in this note. Contract values are reproduced because they are the subject of the analysis; they cannot be attributed to a named party from this document.

## Reproducibility

Every figure derives from read-only SQL against the production call store. No records were modified. The pattern is: resolve each agent's primary outcome key from its live configuration; count validated post-call results matching that key inside the window; count dials, 25-second-plus connections and summed talk duration from the same window's call records; then divide the apportioned fee by the outcome count.

## Suggested citation

Callable AI (2026). *What does one successful outcome cost on an AI voice agent? Seven days of continuous production calling across three anonymised outbound tenants.* Research Note CA-2026-02, 17 August 2026.

---

Prepared by Callable AI from first-party production call data covering 11 – 17 August 2026 (Australia/Brisbane). Amounts in Australian dollars. Tenant identities are withheld by agreement.