

After-call admin, deleted: a voice memo becomes a logged, researched, priced opportunity

The rep talks for sixty seconds. Under thirty seconds later, the opportunity is classified, structured, deduplicated, researched, and staged in the pipeline — behind a rep-approval gate, not blind automation.

60 sec

Replaces 20–30 min of admin per new lead

< 30 sec

Voice memo → staged pipeline entry

3–5 hrs/wk

Returned to selling — 20–30 min × ~10 calls/week

< 2 min

Cited prospect research per Company Card

THE PROBLEM

A telecom agency's field sales team was losing hours every week to after-call admin: typing notes into spreadsheets, building follow-up tasks by hand, researching each prospect manually — and using memory as the deduplication system. A missed follow-up is the quietest way to lose a deal: nothing breaks, nobody complains, the pipeline just shrinks. And every new lead cost 20–30 minutes of work that produced zero new revenue.

WHAT WE BUILT

- **Capture** — transcribes and classifies each memo (new deal, follow-up, meeting, quick note) and extracts the key details into a structured pipeline record.
- **Follow-through** — creates a follow-up project with tasks and files a formatted record in the team's note system; the pipeline is cross-checked so duplicates surface instantly.
- **Research** — an AI Company Card per prospect: industry, HQ, locations, news, buying signals — sources cited, unverified items flagged.
- **Pricing** — at quote time, product mentions become real catalog pricing and deal-level economics.
- **Background agents** — a daily digest, weekly summary, stage-transition automation, and follow-up detection run on schedule, unattended.

The safeguard. Rep-approval gate, not blind automation: the system drafts, structures, and stages; the rep commits what enters the pipeline. Research is cited; unverified items are flagged, never silently filled in.

THE RESULT


Every call is captured the moment it ends — nothing forgotten, nothing re-typed. Prospect intelligence that used to mean manual searching arrives in under two minutes. And the pipeline has total recall: “the landscaping company in Boston with 15 locations” lives in the notes, not in someone's memory.

Does your team's pipeline live in voice mail, memory, and good intentions?

Start with an AI Workflow Teardown →

HOW IT WORKS — RUNTIME FLOW

A voice memo recorded after a call is transcribed, classified, and extracted into a schema-validated pipeline record; research and deduplication run automatically, and everything stages behind a rep-approval gate before it commits. Scheduled background agents keep the pipeline honest without babysitting.



The system drafts, structures, and stages — the rep commits. The human sits where judgment lives, not where the toil lives.

Two design choices stand out. **Schema-validated extraction:** model output must conform to the pipeline schema — malformed records are rejected, not massaged. **Cited research, flagged gaps:** every Company Card claim carries a source; unverified items are marked, never silently filled in.

RELIABILITY AND GUARDRAILS

MECHANISM	WHAT IT DOES
Rep-approval gate	Nothing commits to the pipeline without the rep's sign-off — the human sits where judgment lives, not where toil lives
Cited research, flagged gaps	Every Company Card claim carries a source; unverified items are marked, not filled in
Schema-validated extraction	Model output must conform to the pipeline schema; malformed records are rejected, not massaged
Dependency-disciplined Python	Lightweight build with a deliberately small dependency surface — less to break, less to patch
100+ automated tests, built test-first	Test-driven development throughout, with independent code review, version control, and documented releases
Unattended operation	Background agents run on schedule daily — reliable, repeatable, no babysitting required

INTEGRATION STACK

LAYER	TOOL	ROLE
AI engine	Claude	Structured data extraction, automated web research, schema-validated output
Speech	Speech-to-text transcription pipeline	Voice memo → text
Automation	Python (dependency-disciplined)	Classification, orchestration, background agents
Data & workflow	Excel pipeline, note system, task manager	Where the team already works — zero new tools to adopt
Scheduling	Scheduled background agents	Daily digest, weekly summary, stage transitions, follow-up detection

The same pattern adapts to any team that runs on calls: telecom, insurance, real estate, field services, B2B distribution.

OUTCOMES & METRICS

METRIC	RESULT
Admin per new lead	20–30 min manual work → a 60-second voice memo
Memo → staged pipeline entry	Under 30 seconds of system processing
Weekly selling time recovered	3–5 hours — derived: 20–30 min × ~10 calls/week (arithmetic shown, not asserted)
Prospect research per lead	Manual searching before each conversation → a cited Company Card in under 2 minutes
Missed follow-ups	Recurring (“details slipped through the cracks”) → client-reported: eliminated
Engineering rigor	100+ automated tests (TDD), independent code review, version control, documented releases

Baselines are rep-reported pre-system figures; the weekly figure is derived arithmetic, shown deliberately rather than asserted as a measurement.

TRADEOFFS AND LIMITS

Excel as the pipeline store. Matches how the team already works — zero adoption friction, zero new tooling. A database-backed or CRM-synced store is the natural v2 when scale demands it.

The Company Card flags unverified items rather than suppressing them. Reviewer judgment is preserved on exactly the rows that need it — a design choice, and a real limit on full automation.

Single-team deployment. The build matches the current footprint; multi-team rollout is scoped work, not a toggle.

CONSULTING TAKEAWAY

The unit of adoption is the rep's habit, not the tool. This system works for one structural reason: the input demands nothing new, and the output demands trust. Reps already talk — a 60-second voice memo is zero behavior change. And reps already guard their pipeline — so nothing enters it without their approval. Automation that asks people to change how they work gets abandoned by week three. Automation that removes work while leaving control untouched gets adopted by lunch.

The design lesson is where the gate went: the human sits where judgment lives — committing a record — not where the toil lives — typing it. Get that placement backwards and you've built either a rubber stamp or a data-entry job with extra steps.

Best place to start: an AI Workflow Teardown.

One workflow mapped, scored, and risk-tiered, with a build / no-build recommendation.

Book a 30-min fit call
pete@pdinsights.ai