# How Anthropic's business development team uses Claude to run inbound and outbound at scale

Claude

## Five hours a day, given back

# The time goes into strategic work and customer problems.

| BEFORE | | NOW |
|---|---|---|
| Five hours a day on the inbox | → | **Replies drafted every hour** |
| Accounts researched one at a time | → | **100+ researched overnight** |
| Data requests queued behind a data team | → | **Answered with a prompt** |

Source: the original post · Anthropic Blog, 7 Aug 2026



Amar Harolikar
Decision Sciences & Applied AI

## WHERE THIS STARTED

# "I would spend around 5 hours per day manually responding to inbound interest from prospects .."

### EARLY CAREER, BEFORE ANTHROPIC

- → Lists of hundreds of accounts
- → Investigate each company
- → Find the right contacts
- → Hunt down email addresses

### FIRST MONTHS AT ANTHROPIC

- → Around five hours a day on the sales inbox

**SKILLS AND SCHEDULED TASKS IN CLAUDE COWORK**

### WHAT THE REP OPENS NOW

- → Customer emails already drafted
- → Outbound research already compiled
- → Leads with a first touch written

"A lot of that work is now set up as **skills and scheduled tasks** in Claude Cowork." All of it runs on connections Claude already has: Gmail, Google Calendar, Salesforce, Gong, Apollo, Common Room and the team's data warehouse.

# The document behind every reply

WHAT IT IS

### One sales knowledge base

The questions the sales inbox receives most often, collected in a single document alongside the team's best answers to them.

HOW IT WAS BUILT

### Claude helped build it

The rep pointed Claude at the relevant sources of information rather than writing the document out by hand.

STAYING CURRENT

### Stale answers get flagged

Claude continuously checks the document and flags information that might be stale, which users can validate.

**WHY IT COMES FIRST**

Claude reads this document before drafting any reply the team sends, so it is the source for product facts across the inbound workflow.

**“Build the knowledge base before the workflows.”**

# "It scans a rep's inbox, finds every thread that the rep needs to answer, and drafts a reply."

**1. A THIN SYSTEM PROMPT**

Short instructions, with the working knowledge kept outside the prompt.

**2. THE KNOWLEDGE BASE AS CONTEXT**

The shared question-and-answer document is the source for product facts.

**3. A VOICE PROFILE PER REP**

Built by a voice skill that reads documents, messages and emails the rep has already written, so drafts arrive sounding like the sender.

↓

**RUNS EVERY HOUR**

Drafted replies wait in the inbox for the rep to read, edit and send.

"Claude can generate drafts, but we still **read, edit, and send** them."

# Three more skills the rep built

### Watches for no-shows

Watches Gmail and Google Calendar and notifies the rep when a meeting is missed or a prospect goes dark, so the follow-up happens quickly.

### Drafts the first touch on new leads

Uses the CRM connector to scan for new leads and draft a personalized first message. Runs through the day so leads are not left waiting.

### Pipeline scanner, for keeping Salesforce current

Reads the team's internal guidance on opportunity stages and checks it against what is actually happening in Gmail and Gong. If a customer meeting has moved on to pricing questions, the opportunity should probably progress a stage. Each proposed Salesforce update arrives with the evidence behind it and waits for approval.

### THE CORRECTION GETS RECORDED

When the rep edits or rejects a proposal, the skill records the reason why, so it does not repeat the mistake.

Each of the three ends in something the rep sees: **a notification, a draft, or a proposal with its evidence.**

# “On average, I work upwards of a hundred accounts at any given time.”

## WHAT THE SKILL READS

- → Salesforce
- → Gong
- → Outbound guidance
- → Apollo and Common Room
- → The data warehouse
- → Curated ICP criteria

OVERNIGHT SCHEDULED RUN, ACROSS THE WHOLE BOOK ↓

## WAITING IN THE MORNING, PER ACCOUNT

- → A brief on the account
- → An outbound play
- → A score

## WHAT THE RUN WORKS OUT

Who the team is already in touch with at each account, how that account uses Claude today, and which signals are relevant, all checked against the team's own outbound guidance and ICP criteria before the rep sees it.

Feedback from each rep goes back into the skill. A small memory file and a ledger keep it from repeating work it has already done.

# Scoring the discovery call

**Discovery calls are the part of the outbound motion the team says it is still working to improve. The skill reads Gong transcripts against the discovery call playbook and builds a scorecard for every call.**

**Top three** things done well on the call

**Top three** areas to improve

An explicit **pass or fail** against the team's criteria

One **highest-leverage thing** to practise next

📞 **BDR Call Evaluation**

**Call:** Discovery call with **Customer A** (founder/CEO) — Aug 4, 2026, ~18 min — *Gong link withheld for anonymized output*

**Detected segment: Startup** — seed-stage company ($8M raised) building agents on the API, technical founder as the sole contact, conversation centered on model performance, credits, and scaling production workloads.

**Overall score:** 4.0 / 5

| Dimension | Score | One-line summary |
|---|---|---|
| Call Opening & Prep Signals | 3/5 | Warm rapport and clean agenda-set, but joined 4 min late and no research hypothesis |
| 70/30 Rule | 5/5 | Prospect 70.6% / you 29.4% (computed from raw transcript word counts) — textbook |
| Strategic Deep Dive | 4/5 | Strong question ladder (use case → moat → GTM → competition), but the numbers went unprobed |
| Next Step & AE | 4/5 | Right call to keep it BDR-owned; concrete, time-bound follow-ups |

The call-coach scorecard, from the original post. Demo data, anonymized by Anthropic for publication.

# Three things that used to queue behind a data team

REQUEST 01

### Usage trends for a top account

An account executive asks, and a legible descriptive dashboard is a prompt away.

REQUEST 02

### Undiscovered usage sweep

Runs across an AE's full book and returns every account already using the product where no sales opportunity exists yet.

REQUEST 03

### Who to invite to the webinar

No skill existed for this one. Claude checked usage data and CRM history across the book, scored each account against the ICP, and flagged the best fits with the contacts worth inviting.

No skill existed for the webinar request, and **a prompt was enough**. The other two are described as prompts as well, run against usage data and CRM history.

# Advice for business development teams

01

## Knowledge base first

Collect the questions your team answers repeatedly, and your best answers, into one external-facing document. Point Claude at your product docs and team channels to build the first version.

02

## Show it how you work

Messages that worked, your ideal customer profile, and each rep's writing style, so drafts arrive sounding like the person sending them.

03

## A person on every send

Claude generates the draft. The rep reads it, edits it, and sends it.

04

## Share the skills

The team keeps its most-used skills in a shared plugin, promoting one there once reps use it consistently in daily work.

05

## Keep them general

Segments, books and workflows differ across reps, so a shared skill adapts rather than being scoped to one person's routine.

06

## Write feedback back in

When you dismiss a hook or correct a draft, have Claude record the reason in the skill so the same mistake does not repeat.

**“My best advice? Just start experimenting.”**

# My read

MY OWN VIEW. NOT FROM THE ANTHROPIC POST.

I have spent years on campaign work, mostly outbound and mostly retail banking. My work now is mostly app development, automation and internal tooling, and the same few things decide whether any of it works.

**The database first.** Accuracy, validation and access control, because everything built on top inherits whatever is wrong down there.

**Then the skills.** Anything repeatable goes into one, and a skill is really a record of how we do a thing and how we want it done. The AI does not need me to hand it Python, it needs my context, and that sits in the skills and in the knowledge documents next to them.

**Then the workflows.** Crons and scheduled runs on a server or a local machine, producing triggers, alerts, trackers and scanners, so the work happens without anyone remembering to start it.

**And a person still reviews.** The pipeline runs on its own, but there is a check at every step that counts. Validation reports come out, I read them, and if something does not look right I go and dig into it. AI or no AI, the buck stops with me.

# Source

**"How Anthropic's business development team uses Claude to run inbound and outbound at scale"** · John Albert, business development at Anthropic. Published 7 August 2026. UI mockups in the original use synthetic data. These pages are a scannable summary of it.

Live macro data, database AI, quants and MCP servers, plus 45 live apps and 200+ build guides: tigzig.com.