

# Account Signal Agent

An AI agent that researches target accounts and returns what to say to each one — and what not to. Built in an afternoon, with no prior coding, automation, or API experience.

**Parth Vashishta**    B2B Marketing — ABM & Demand Generation    [parthvashishta1995@gmail.com](mailto:parthvashishta1995@gmail.com)

| BUILD TIME | API SPEND        | PER ACCOUNT | RUNTIME        | PRIOR EXPERIENCE |
|------------|------------------|-------------|----------------|------------------|
| 4–5 hours  | \$2.36 of \$5.00 | \$0.36      | 90 sec/account | None             |

## 01 The problem

Account-based marketing starts with a defined account list. Everything that happens next is the bottleneck.

Researching one account properly — what they actually do, who would sit on the buying committee, whether this is the right moment to reach out, what would land and what would backfire — takes thirty to forty-five minutes if you do it honestly. Multiply that across forty target accounts in a quarter and it becomes most of a week.

So it doesn't happen. The top five accounts get researched properly. The rest get a LinkedIn glance and a templated sequence, and then everyone wonders why reply rates are what they are.

This agent doesn't replace that judgment. It does the first pass on every account, so the judgment starts from something rather than from a blank page.

70% quality on all forty accounts beats 100% quality on the three there was time for.

## 02 What it produces

Input is a Google Sheet with two columns: company name and domain. Output is a row per account, written back to a second tab, with six fields.

| FIELD       | WHAT IT ANSWERS                                                                                                                                                                                        |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| fit         | Is this account genuinely in our ICP? Rated Strong, Possible or Poor, with the reasoning stated.                                                                                                       |
| committee   | Three to five roles likely involved in the purchase, each labelled economic buyer, technical evaluator, end user, or champion — with seniority calibrated to that company's actual size and structure. |
| trigger     | The single strongest reason to reach out now, with its date. Or the words "No trigger found."                                                                                                          |
| hook        | The specific problem to open with, referencing something concrete about that account.                                                                                                                  |
| proof_point | Which of our proof points to lead with here, and why that one for this account.                                                                                                                        |
| avoid       | What not to say to this account, given what the research revealed.                                                                                                                                     |

The last field is the one that justifies the build. Plenty of tools will summarise a company. Very few will tell you which perfectly reasonable-sounding email would get you blocked.

A real result

Verbatim output from a test run against Hootsuite, with a hypothetical omnichannel helpdesk product configured as the thing being sold:

Hootsuite

FIT: STRONG

TRIGGER

CEO transition on April 13, 2026, with returning founder Ryan Holmes stating a new priority of "investing in AI to help both our customers and ourselves move faster and work smarter" — a fresh, stated mandate that opens a window for tooling conversations.

HOOK

Hootsuite's support org just went through a 20% workforce reduction (Oct 2025) while facing heightened public complaint volume from the ICE/DHS backlash (Jan–Feb 2026) — fewer agents are likely handling more escalations across separate social, email and phone systems, with no single contextual record of each contact.

AVOID

Don't reference or comment on the DHS/ICE contract, the backlash, or the layoffs as talking points — these are sensitive, unresolved and reputationally charged; also avoid assuming a huge enterprise-scale support org (only 30+ staff confirmed) or claiming known complaints about response times, since the brief states none were found.

Nobody told the agent about the ICE contract, the layoffs, or the CEO change. It found all three, worked out which was the opening and which was the landmine, and said so. The `fit` rating was reached by matching the ICP against a job advert that listed the company's actual contact-centre stack — IVR, ACD, ticketing, chat and WFM — which is also what it selected the proof point on.


It was not uniformly generous. Two of the three test accounts came back "Possible" rather than "Strong," each with a specific reason — one because its end users are developers reached through a community forum and GitHub, which the agent noted is "a different dynamic than agent-to-customer ticketing." That is an expensive thing to discover on a call instead.

|   | A         | B                                                                                                                                                                                                                                                                                                                                                                                         | C                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | D                                                                                                                                                                                                                                                                                       | E                                                                                                                                                                                                                                                                                                                                                                    | F                                                                                                                                                                                                                                                                                                   | G                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1 | company   | fit                                                                                                                                                                                                                                                                                                                                                                                       | committee                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | trigger                                                                                                                                                                                                                                                                                 | hook                                                                                                                                                                                                                                                                                                                                                                 | proof_point                                                                                                                                                                                                                                                                                         | avoid                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| 2 | Miro      | Possible - Miro has the scale (1,800-2,800+ employees, large enterprise base) but the brief shows support is currently ticket-and-community based rather than already fragmented across 3-5 live channels, so it's a fit on complexity and pain but not yet as clear match to the "multichannel fragmentation" ICP trigger.                                                               | ECONOMIC BUYER, VP/Head of Customer Success (Americas) - owns the post-sale org budget and is actively building out the team TECHNICAL EVALUATOR, Support Operations/Support Representative last handling SSO and API troubleshooting - would assess integration with existing ticket workflows END USER, Support Representatives/agents feeling technical and billing tickets CHAMPION, Support Manager dealing directly with escalation delays and customer complaints - has clearest incentive to fix tooling                                                                                                                                                                                                                                               | Miro's back-to-back acquisitions of Butter (Jan 2025) and Refrigo (announced March 24, 2025) are adding product surface area and support complexity right as the brief shows no parallel investment in support infrastructure.                                                          | Miro's support today is essentially ticket-and-community based - the chat is locked behind Enterprise plus Premium Support, and there's no phone or social ticketing - while reviewers describe unanswered complaints and slow escalations on billing and technical issues.                                                                                          | Cutting average first response time from 8 hours to under 1 - this directly counters the recurring reviewer complaints about Miro support going unanswered or resolving too slowly.                                                                                                                 | Don't cite Miro's user/customer figures precisely (SSO vs SGM vs 2020 conflict in the brief), don't claim they already run true omnichannel support, don't pitch this as a sales-team tool since the buying centre sits in CX/support, and don't frame the layoffs as evidence of an understaffed support team since the brief doesn't confirm layoffs to that org specifically.                                                                                                                                                                                                                                           |
| 3 | Postman   | Possible - Postman meets the headcount and high-volume criteria (500+ employees, 40M+ developers, 500K+ organizations) and shows signs of multichannel support (phone, ticketing, community/DisCORD, GitHub), but the brief has no confirmed evidence of 30+ support agents or three enterprise-grade channels operating as fragmented tools, so fit is plausible but not fully verified. | Technical Support Manager (building out EMEA/APAC team) - ECONOMIC BUYER, likely owns the support tooling budget and headcount planning / Director/Head of Customer Support (implied owner of the Support Center, unnamed in brief) - ECONOMIC BUYER or CHAMPION, Technical Support Engineer handles phone/webinar/share support, GitHub issues, community - END USER, IT/Security or Trust Center owner - TECHNICAL EVALUATOR, would vet any new platform against Postman's compliance/security posture / Community Manager (DisCORD community, 15M+ developers) - END USER, stakeholder for any channel unification affecting community workflows                                                                                                            | March 2026 launch of Postman's rebuilt AI-native platform, layered on top of three recent acquisitions (Stable, Fern, and earlier deals) - this kind of product expansion typically spikes ticket volume and complexity right when the support org is also being rebuilt after layoffs. | Postman's support already spans phone, a ticket-based Support Center, DisCORD community, and GitHub issue tracking with no evidence of a unified record connecting them - and they're now expanding into EMEA/APAC, while absorbing three acquisitions and a major platform relaunch, which is exactly when conversations start fragmenting and agents lose context. | Replaced multiple disconnected tools with one omnichannel workspace - Postman's support already runs across at least four separate surfaces (Support Center, DisCORD, GitHub, phone) with no sign of a unifying system, making tool consolidation the most concrete and relevant pain to lead with. | Don't pitch this as a sales-motion tool or reference their sales phone line - the buying centre here is support/CX, not sales. Don't cite the unverified 3.5x headcount or the 82% layoff figure as confirmed first since the brief flags both as unverified. Don't claim they have response-time problems or a failing support stack - no public complaints or migration evidence exists, so that would be fabricating pain. Avoid treating them like a standard enterprise B2B support org - their end users are largely developers via community/GitHub, which is a different dynamic than agent-to-customer ticketing. |
| 4 | Hootsuite | Strong - Hootsuite has 900-1800 employees (well above the 500+ threshold), a structured support org of 30+ agents spanning at least four channels (social/Twitter/X, Facebook, email, phone), and job postings explicitly citing separate IVR, ACD, ticketing, chat and WFM systems, which is precisely the fragmented-tooling pattern our platform is built for.                         | ECONOMIC BUYER, VP, Professional Services - the only VP-level role identified in the brief with cross-functional, multi-location scope, plausible final sign-off for a tooling change of this size TECHNICAL EVALUATOR, Manager, Customer Support - oversees 30+ operational staff and is required to know IVR, ACD, ticketing, chat and WFM systems, making them the natural evolution of a consolidated platform CHAMPION, Team Lead, Enterprise Customer Support - owns support delivery across email, phone, chat/messaging, social and public forums, and directly feels the pain of fragmentation day to day END USER, Customer Support Repertoire - frontlines agent handling tickets across the channel mix, whose workflow would change most directly | CEO transition on April 13, 2026, with returning founder Ryan Holmes stating a new priority of investing in AI to help both our customers and ourselves move faster and work smarter - a fresh, stated mandate that opens a window for tooling conversations.                           | Hootsuite's support org just went through a 20% workforce reduction (Oct 2025) while facing heightened public complaint volume from the ICE/DHS backlash (Jan-Feb 2026) - fewer agents are likely handling more escalations across separate social, email and phone systems, with no single contextual record of each contact.                                       | Replaced multiple disconnected tools with one omnichannel workspace - Hootsuite's own job postings list separate IVR, ACD, ticketing, chat and WFM systems as required experience, which is direct evidence of the exact tool sprawl this proof point addresses.                                    | Don't reference or comment on the DHS/ICE contract, the backlash, or the layoffs as talking points - these are sensitive, unresolved and reputationally charged, also avoid assuming a huge enterprise-scale support org (only 30+ staff confirmed) or claiming known complaints about response times, since the brief states none were found.                                                                                                                                                                                                                                                                             |

The full output tab. Three accounts, seven columns, written back automatically. The `committee` column labels each role as economic buyer, technical evaluator, end user or champion, with the reasoning for each attached.

03 How it works

Ten steps in a single workflow. Two calls to Claude per account: one to gather, one to judge.



One phase decides its own actions at runtime. The other three run a fixed recipe. That distinction is the whole architecture, and it is also where 89% of the cost sits.

What makes one step an agent

In the research phase, Claude is given a brief and access to web search, and nothing else. It writes its own queries, reads what returns, decides whether that was enough, and searches again if it wasn't. On a single account it independently chose to search for headcount, then support channels, then acquisitions and layoffs, then support job listings, then a specific geographic expansion it had picked up in the previous result. None of those queries were written by a human.

The other three phases are deliberately not agentic. Extracting text, turning a brief into six fields, appending a row — these have known inputs and known outputs. Handing control to a model there would add cost, latency and variance in exchange for nothing.

Agency is a cost, not a virtue. It buys the ability to handle open-ended input, and it is paid for in reproducibility, cost predictability and debuggability. Spend it where the task is genuinely open-ended and nowhere else.

04 Five decisions

REUSABILITY

The product being sold lives in a config node, not in the prompts

Product name, category, description, ICP and proof points sit in five fields at the top of the workflow. Every prompt downstream reads from them. Nothing about the product is hardcoded anywhere else.

This means re-pointing the agent at a different product, category or ICP is a five-field edit rather than a prompt rewrite — which is what makes it a tool rather than a one-run demo. The alternative was to hold config in a spreadsheet row, which would let a non-technical colleague change it without opening the workflow; that is the right call for a team tool and the wrong one for a solo build.

DATA SOURCING

Claude's own web search, not a third-party search API

The alternatives were Tavily, Brave or SerpAPI on their free tiers, or scraping. Each returns a list of links and snippets, which means building the loop yourself: search, evaluate, refine the query, search again, then fetch and parse the pages. That loop is the most fragile part of any research agent.

Claude's server-side web search runs that loop inside a single API call and returns cited sources. It costs a cent per search rather than nothing, which at this volume is roughly sixty pence across the entire build. Paying pennies to delete the most breakable component was straightforwardly correct.

## PIPELINE SHAPE

### Two Claude calls, not one

Gathering and judging were kept separate. The research call's objective is recall — find everything, filter nothing. The analysis call's objective is judgment, working only from facts already committed to.

Merging them is measurably worse: when a model writes the narrative and sources the facts in the same breath, the narrative starts pulling the facts around. It also destroys debuggability — with two calls, a bad row can be traced to either bad research or bad interpretation, and they are fixed in completely different places.

## RELIABILITY

### Designed against fabrication from the first prompt

Four specific guards, each closing a known failure:

- **The domain is declared authoritative.** Search results for "Miro" returned Miro Forestry, Mircom and Mirova. The brief opened by naming and excluding all three.
- **"Not found" is an explicitly permitted answer**, stated twice. Ask a model for buying signals and it will produce buying signals, because that is what was requested. Making absence a valid answer is what separates trustworthy output from plausible output.
- **Today's date is injected**, and every claim requires a date. Without this a 2024 funding round gets described as recent.
- **Proof points are supplied, never generated.** The model's job is to select from three given claims, not to invent a customer result. A fabricated case study is the single most dangerous thing this agent could produce, because it is the one an SDR would paste straight into an email.

That last guard visibly worked. For one account the model refused to use a response-time statistic, writing: *"no public complaints or migration evidence exists, so that would be fabricating pain."* For a different account it led with that same statistic, because that account's research contained reviewers complaining about unanswered tickets. Same three claims, different selection, reasoning stated each time.

## MODEL ROUTING

### The best model on the step that looks like grunt work

Conventional routing puts a cheap fast model on gathering and an expensive one on writing. That is backwards here. Research looks mechanical, but the model is composing its own search queries, and query formulation is where quality is actually decided. A weaker model searches the company name and stops.

Bad research poisons everything downstream and no amount of clever prompting recovers from it. So the strong model goes at the top of the funnel, where its output becomes everyone else's input.

## 05 What broke

Four things went wrong. Each was more instructive than the parts that worked.

FAILURE 01

The purpose-built node couldn't do the job

n8n ships a native Anthropic node with a Tools panel, which looked like the obvious way to enable web search. It isn't there — the panel offers n8n's own tool types only. Further digging turned up an open bug report on the alternative path: the node performs the search, gets billed, and returns encrypted content it never converts back to readable text, because it completes the first half of the tool exchange and stops.

The fix was to drop the friendly node and call the Anthropic API directly with a raw HTTP request, hand-writing the JSON body including the tool definition.

**Lesson:** the abstraction that makes a thing easy is also what hides the thing from you. Being forced one layer down turned out to be an upgrade — the raw call is now the part of this build that can be explained in the most detail.

FAILURE 02

Silent truncation, twice

The output-length cap was set too low. The research briefs stopped mid-sentence, losing exactly the two sections that mattered — recent developments and buying signals. It happened again on the analysis step for a different reason: extended reasoning was consuming the entire token budget before the model reached its answer, so two of three accounts returned pure reasoning and no output at all.

**Lesson:** nothing errors when this happens. The API returns success, the workflow goes green, and the writing simply stops. Any failure that presents as success is worth more attention than one that presents as a crash.

FAILURE 03

Unspecified format means improvised format

The prompt asked for prose "under four headings" without saying what the headings should look like. Three accounts came back in three different formats — one numbered with markdown hashes, one in bold caps, one with an invented title line on top. Fixed by specifying the four headings verbatim and forbidding a preamble.

**Lesson:** anything left unspecified is decided by the model, freshly, every single time. Consistency is not a default you can rely on — it is an instruction you have to write.

FAILURE 04

The cost estimate was wrong by a factor of three

Budgeting assumed the expensive part of a search-enabled call was what the model writes. It isn't. Search results pull page content *into* the model's context — roughly 99,000 input tokens per account, against about 2,500 written out. Input was 73% of the cost of every research call.

**Lesson:** for any agent with tools, model the cost of what comes back in, not what goes out. The correction also produced the single most valuable habit in the build: freezing expensive outputs in place so that hours of downstream iteration cost nothing.

06 Cost

Total spend across the entire build, including every failed attempt, every re-run and all prompt iteration:

| COMPONENT         | BASIS                                               | PER ACCOUNT |
|-------------------|-----------------------------------------------------|-------------|
| Research — input  | ~99k tokens of search results returned into context | \$0.20      |
| Research — output | ~2.5k tokens, a 500-word sourced brief              | \$0.03      |

| COMPONENT                 | BASIS                               | PER ACCOUNT |
|---------------------------|-------------------------------------|-------------|
| Web searches              | 3–5 queries, chosen by the model    | \$0.05      |
| Analysis                  | Brief in, six structured fields out | \$0.04      |
| n8n, Google Sheets        | Free tier and trial                 | \$0.00      |
| Marginal cost per account |                                     | \$0.32      |

Against a human benchmark of thirty to forty-five minutes per account, the agent runs in about ninety seconds. Development consumed \$2.36 of a \$5.00 balance, which included re-running the research phase twice after the truncation was found.

At scale the economics stay reasonable and the constraint moves: a hundred accounts costs roughly \$32, but the search volume becomes the lever worth tuning, not the model.

## 07 Limits

Stated plainly, because a portfolio piece that claims no weaknesses is not describing something real.

- **It only sees what is public.** Buying signals that live inside a company are invisible to it. This also means the choice of product category determines whether the agent finds anything at all — sell something whose triggers never surface publicly and it will correctly return "no trigger found" for most of the list while appearing broken.
- **A failure late in a run costs the whole run.** All accounts move through each step together, so a permanent error on account six loses one through five as well. The fix is a loop that pushes each account through the full chain independently and writes its row before starting the next.
- **Nothing checks the output against the research.** A cheap second pass verifying every claim traces back to a source in the brief would close the last gap on fabrication. It was cut for budget, not because it wasn't warranted.
- **No firmographic data source.** Headcount figures come from whatever third-party trackers surface in search, and they disagree — one account returned estimates ranging from 1,400 to 2,840 employees. The agent flagged the discrepancy rather than picking one, which is the right behaviour, but a real data provider would remove the ambiguity.
- **Re-runs duplicate rather than update.** There is no change detection, so running the same list twice appends a second set of rows rather than refreshing the first.

## 08 Where else this applies

The pattern generalises further than the original use case. Everything below is a change of configuration and prompt, not of architecture.

| USE CASE                  | WHAT CHANGES                                                                                                                       |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------|
| Pre-call briefs for sales | Trigger on a booked meeting rather than a manual run; output to CRM instead of a sheet.                                            |
| Territory prioritisation  | Run the full territory and sort by the fit rating to decide where the quarter's effort goes.                                       |
| Competitive monitoring    | Point the account list at competitors and change the six output fields to positioning shifts, pricing changes and hiring patterns. |
| Event target lists        | Feed in a conference attendee list; return which accounts are worth booking time with and what to open with.                       |
| Renewal risk scanning     | Run against the existing customer base, looking for leadership changes, layoffs and competitor announcements.                      |
| Investment screening      | Same research engine, different judgment prompt — the gathering step is entirely domain-agnostic.                                  |

## 09 Stack

|                                                                                                                             |                                                                                                    |                                                                                          |
|-----------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------|
| ORCHESTRATION                                                                                                               | MODEL                                                                                              | RESEARCH                                                                                 |
| n8n Cloud<br>14-day trial; workflow exported as JSON so it survives the trial and moves to a self-hosted instance unchanged | Claude Sonnet 5<br>Anthropic API, both calls                                                       | Anthropic web search tool<br>Server-side; the model issues and evaluates its own queries |
| DATA IN / OUT                                                                                                               | INTEGRATION                                                                                        | CODE WRITTEN                                                                             |
| Google Sheets<br>Two tabs: account list in, structured research out                                                         | Direct HTTP to the Messages API<br>Hand-written JSON body including the web search tool definition | None<br>Two one-line expressions to parse API responses                                  |

## 10 What this was for

Two things. The first was to have a working tool — this genuinely does a job worth doing, and it does it for a third of a pound an account.

The second was to be able to explain and extend it under pressure. That constraint shaped the build more than any other. Every design choice above was made deliberately and can be defended, reversed, or rebuilt from scratch — including the ones that turned out to be wrong.

Point it at a different product and it researches accounts for that instead. That takes about a minute.

## A Full output, three accounts

Everything the agent returned, unedited. The `committee` column has been broken onto separate lines for readability; nothing else has been changed, cut or reordered. The product configured as the thing being sold was a hypothetical omnichannel helpdesk platform, ICP set to B2B SaaS companies of 500+ employees with 20+ support agents across three or more channels.

Miro

FIT: POSSIBLE

FIT REASONING

Miro has the scale (1,800–2,800+ employees, large enterprise base) but the brief shows support is currently ticket-and-community based rather than already fragmented across 3+ live channels, so it's a fit on complexity and pain but not yet a clean match to the "multichannel fragmentation" ICP trigger.

BUYING COMMITTEE

ECONOMIC BUYER

VP/Head of Customer Success (Americas) — owns the post-sale org budget and is actively building out the team

TECHNICAL EVALUATOR

Support Operations / Support Representative lead handling SSO and API troubleshooting — would assess integration with existing ticket workflows

END USER

Support Representatives / agents fielding technical and billing tickets

CHAMPION

Support Manager dealing directly with escalation delays and customer complaints — has clearest incentive to fix tooling

TRIGGER

Miro's back-to-back acquisitions of Butter (Jan 2026) and Reforge (announced March 24, 2026) are adding product surface area and support complexity right as the brief shows no parallel investment in support infrastructure.

## HOOK

Miro's support today is essentially ticket-and-community based — live chat is locked behind Enterprise plus Premium Support, and there's no phone or social ticketing — while reviewers describe unanswered complaints and slow escalations on billing and technical issues.

## PROOF POINT

Cutting average first response time from 8 hours to under 1 — this directly counters the recurring reviewer complaints about Miro support going unanswered or resolving too slowly.

## AVOID

Don't cite Miro's user/customer figures precisely (60M vs 90M vs 200K conflict in the brief), don't claim they already run true omnichannel support, don't pitch this as a sales-team tool since the buying centre sits in CX/Support, and don't frame the layoffs as evidence of an understaffed support team since the brief doesn't confirm layoffs hit that org specifically.

# Postman

FIT: POSSIBLE

## FIT REASONING

Postman meets the headcount and high-volume criteria (500+ employees, 40M+ developers, 500K+ organizations) and shows signs of multi-channel support (phone, ticketing, community/Discord, GitHub), but the brief has no confirmed evidence of 20+ support agents or three enterprise-grade channels operating as fragmented tools, so fit is plausible but not fully verified.

## BUYING COMMITTEE

**ECONOMIC BUYER** Technical Support Manager (building out EMEA/APAC team) — likely owns the support tooling budget and headcount planning

**BUYER / CHAMPION** Director or Head of Customer Support (implied owner of the Support Center, unnamed in brief)

**END USER** Technical Support Engineer (handles phone / screen-share support, GitHub issues, community)

**TECHNICAL EVALUATOR** IT / Security or Trust Center owner — would vet any new platform against Postman's compliance and security posture

**END USER** Community Manager (Discord community, 15M+ developers) — stakeholder for any channel unification affecting community workflows

## TRIGGER

March 2026 launch of Postman's rebuilt AI-native platform, layered on top of three recent acquisitions (liblab, Fern, and earlier deals) — this kind of product expansion typically spikes ticket volume and complexity right when the support org is also being rebuilt after layoffs.

## HOOK

Postman's support already spans phone, a ticket-based Support Center, Discord community, and GitHub issue tracking with no evidence of a unified record connecting them — and they're now expanding into EMEA/APAC while absorbing three acquisitions and a major platform relaunch, which is exactly when conversations start fragmenting and agents lose context.

## PROOF POINT

Replaced multiple disconnected tools with one omnichannel workspace — Postman's support already runs across at least four separate surfaces (Support Center, Discord, GitHub, phone) with no sign of a unifying system, making tool consolidation the most concrete and relevant pain to lead with.

## AVOID

Don't pitch this as a sales-motion tool or reference their sales phone line — the buying centre here is support/CX, not sales. Don't cite the unverified 3.5K headcount or the 62% layoff figure as confirmed fact since the brief flags both as unverified. Don't claim they have response-time problems or a failing support stack — no public complaints or migration evidence exists, so that would be fabricating pain. Avoid treating them like a standard enterprise B2B support org — their end users are largely developers via community/GitHub, which is a different dynamic than agent-to-customer ticketing.

## Hootsuite

FIT: STRONG

### FIT REASONING

Hootsuite has 900–1800 employees (well above the 500+ threshold), a structured support org of 30+ agents spanning at least four channels (social/Twitter/X, Facebook, email, phone), and job postings explicitly citing separate IVR, ACD, ticketing, chat and WFM systems, which is precisely the fragmented-tooling pattern our platform is built for.

### BUYING COMMITTEE

**ECONOMIC BUYER** VP, Professional Services — the only VP-level role identified in the brief with cross-functional, multi-location scope, plausible final sign-off for a tooling change of this size

**TECHNICAL EVALUATOR** Manager, Customer Support — oversees 30+ operational staff and is required to know IVR, ACD, ticketing, chat and WFM systems, making them the natural evaluator of a consolidated platform

**CHAMPION** Team Lead, Enterprise Customer Support — owns support delivery across email, phone, chat/messaging, social and public forums, and directly feels the pain of fragmentation day to day

**END USER** Customer Support Advocate — frontline agent handling tickets across the channel mix, whose workflow would change most directly

### TRIGGER

CEO transition on April 13, 2026, with returning founder Ryan Holmes stating a new priority of "investing in AI to help both our customers and ourselves move faster and work smarter" — a fresh, stated mandate that opens a window for tooling conversations.

### HOOK

Hootsuite's support org just went through a 20% workforce reduction (Oct 2025) while facing heightened public complaint volume from the ICE/DHS backlash (Jan–Feb 2026) — fewer agents are likely handling more escalations across separate social, email and phone systems, with no single contextual record of each contact.

### PROOF POINT

Replaced multiple disconnected tools with one omnichannel workspace — Hootsuite's own job postings list separate IVR, ACD, ticketing, chat and WFM systems as required experience, which is direct evidence of the exact tool sprawl this proof point addresses.

## AVOID

Don't reference or comment on the DHS/ICE contract, the backlash, or the layoffs as talking points — these are sensitive, unresolved and reputationally charged; also avoid assuming a huge enterprise-scale support org (only 30+ staff confirmed) or claiming known complaints about response times, since the brief states none were found.

THREE ACCOUNTS · TWO "POSSIBLE", ONE "STRONG" · TWO DIFFERENT PROOF POINTS SELECTED · ONE PROOF POINT EXPLICITLY REFUSED