




The Reply Gap.

Where AI-powered outbound loses revenue: from the first signal to the moment a prospect answers, and what happens next.



Three acts, one question: where does revenue leak?

A discovery-led approach. Start from a seller's real workflow, follow the evidence, and only then shape a solution.

01

Market research

What do the tools promise, and where do their users say they fall short?

- 1.1 Choose the segment
- 1.2 Study competitors
- 1.3 Read review signals
- 1.4 Form a hypothesis

02

User research

Test the hypothesis against independent evidence and find the sharpest pain.

- 2.1 Listen to the segment
- 2.2 Analyze insights
- 2.3 Map persona and journey
- 2.4 Prioritize to P0

03

Solution and proof

Shape the answer, decide how success gets measured, then prototype it.

- 3.1 Shape the solution
- 3.2 Define success
- 3.3 Prototype and PRD

Market research

Six tools, three segments and a lot of reviews, narrowed down to one testable hypothesis.

Segment

Competitors

Review signals

Hypothesis

Three candidates, scored before committing to one

A

Selected

SDRs at SMB and mid-market B2B SaaS

Reach
High, active on G2, Capterra, Reddit

Pain intensity
High, daily quota pressure

Willingness to pay
Medium, needs manager buy-in

Fit with research
Strong, matches the landscape

Score **11**/12

B

Founder-led, early-stage startups

Reach
Medium, scattered, less vocal

Pain intensity
High, time-starved generalists

Willingness to pay
Low, bootstrapped budgets

Fit with research
Partial, a different buyer motion

Score **8**/12

C

RevOps at 200 to 500 person scale-ups

Reach
Medium, smaller pool, B2B forums

Pain intensity
Medium, process pain, not personal


Willingness to pay
High, owns the tooling budget

Fit with research
Partial, the buyer isn't the user

Score **9**/12

Six tools, priced for very different teams

Published monthly prices on a log scale. Per-seat plans multiply with every rep, so the gap is wider than it looks.



Strong fit

Apollo.io
Database, sequencer and dialer

Pain: 65 to 70% data accuracy, 15 to 25% bounces

Stretch fit

AiSDR
Autonomous outreach agent

Pain: Newer, few public reviews

Partial fit

Reply.io
Multichannel sequencing

Pain: Add-ons push cost to \$150 to \$187 a user

Partial fit

Clay
Enrichment and orchestration

Pain: Needs a GTM engineer

Reviewed fit

Instantly
Sending, warmup and replies

Pain: Confusing modular pricing

Stretch fit

Artisan
Full-stack autonomous AI SDR

Pain: Priced for funded teams

What users say each tool does better, and worse

Apollo.io

- + All-in-one prospecting and outreach
- + Seamless contact discovery
- + Native CRM integrations
- Overwhelming at first
- Contact data isn't always accurate
- Credits vanish, unclear costs

AiSDR

- + Signal-based targeting: funding, role changes
- + Proactive success team
- + Saves real headcount vs manual
- Needs daily ICP and reply oversight
- Weak reporting, no A/B testing
- Credits run out mid-campaign

Reply.io

- + Genuine multichannel without an ops team
- + Strong AI plus human support
- + Fast, intuitive setup
- Reply-based removal is unreliable
- Basic reporting, no unified outbox
- Limited branching across channels

Clay

- + Unmatched enrichment depth
- + Its agent verifies almost anything, per lead
- + Plugs into almost any tool
- Steepest learning curve of the six
- Credit pricing scales quickly
- Not a sender or CRM on its own

Instantly

- + Leads, sending, warmup and replies in one
- + Unlimited inboxes, not priced per seat
- + Strong deliverability tooling
- Confusing modular pricing
- Basic reporting, billing disputes
- Replies can drop suddenly

Artisan

- + Deep personalization, 3.7% vs under 1% replies
- + Finds, researches, writes and sends
- + Hands-on success team
- Gap between autonomy promise and execution
- Clunky UI, work doesn't always save
- Can target people who've left

Finding leads works. Turning them into replies doesn't.

WHAT REVIEWS SHOWED

List building works across all six tools. Data drift exists, but it's minor and well known.



WHERE THE GAP SITS

Outreach is where tools diverge. AiSDR and Reply.io lack branching and A/B tests. Instantly has both, blunted by trust issues.



WHAT IT MEANS

The unmet need is a control gap: adaptive sequencing that SDRs can also trust and steer.

THE HYPOTHESIS

As an SDR at a small or mid-market B2B SaaS company, I need outreach sequencing with **adaptive branching** and **built-in A/B testing**, while keeping **meaningful control** over the AI's decisions, so I can turn more well-prospected leads into real conversations and paying customers.

User research

Putting the hypothesis against independent evidence, so the same voices aren't heard twice.

Listen

Insights

Persona

Journey

P0

Four questions, six independent sources

Q1 Where does outreach break once a sequence is live?

Q2 What happens when a reply isn't a clear yes or no?

Q3 What does testing look like at one SDR's volume?

Q4 What would it take to trust AI to steer a sequence?

Bridge Group SDR Metrics 2025

351 B2B companies

Activity, conversations, quota

Salesforce State of Sales

Sales pros worldwide

Selling time, AI adoption, data trust

Gartner buyer survey, Jun 2025

632 B2B buyers

Reactions to irrelevant outreach

Gartner sales AI outlook, Nov 2025

Orgs rolling out AI agents

Whether agents lift productivity

Deepline, RevBoss, Broadn

Operators running outbound

AI SDRs without oversight

Unify, Gmail and Yahoo rules

GTM engineers, inbox providers

Hard limits on testing and sending

| A note on method: this round is secondary research. G2 reviews shaped the Act one hypothesis, so they were deliberately left out here.

Five patterns, in five numbers

01

4.1

quality conversations from 112 daily activities

Every conversation is expensive

Only 60% of SDRs hit quota, the lowest on record.

Bridge Group

02

+40 to 50%

more replies from the first follow-up alone

Replies come late, rarely clean

The value lives in steps 2 to 4, and most replies aren't a yes.

Instantly, Belkins

03

400 to 600

contacts per variant to detect a 5 point lift

A/B testing breaks at SDR volume

Testing has to pool results and judge on meetings.

Unify

04

<40%

of sellers will say AI agents helped by 2028

Autonomy without grounding loses trust

Control is the condition for adoption.

Gartner, Broadn

05

73%

of B2B buyers avoid suppliers who send irrelevant outreach

Bad outreach costs more than none

Exiting a sequence is a branch, not a failure.

Gartner

Validated

Rigid one-path sequences and missing control are real, independent of G2.

Refined

A test button won't work at SDR volume. The need is a learning loop judged on meetings, and the sharpest pain sits at the reply.

Meet Maya, a composite SDR

Built from the research, not a single interviewee.



Maya Brooks, 26

SDR at an 80-person B2B SaaS company
selling into US mid-market

In role: 1.5 years, eyeing the AE seat

Stack: Apollo, HubSpot, LinkedIn

Her day: about 110 touches

Measured on: meetings her AE accepts

"I don't need more leads. I need my sequence to notice when someone actually answered."

GOALS

- Hit her number every month, not just the good ones
- Hours on conversations, not admin
- Prove she can run a book

FRUSTRATIONS

- A "not now" or "talk to Priya" pulls the lead into manual work
- Numbers too small to trust
- AI drafts that sound wrong, found after sending


ON AI

- Happy to hand over research and drafts
- Wants the why before anything goes out
- Auto-send only where she's seen it get it right

HER JOB TO BE DONE


When a prospect finally responds, help me give them the right next step quickly, without rebuilding the sequence by hand or losing control of what goes out in my name.

One sequence through Maya's week, and where it drops



Seven problems, scored out of 20

Frequency, severity, market gap and feasibility, each scored 1 to 5. Gap scores come from Act one.



THE P0 PROBLEM

When a prospect replies with anything other than a clear yes or no, Maya's sequence has no next step. She handles it by hand or the lead goes quiet, and AI tools that try to help don't show her why.

HOW MIGHT WE

Route every reply to the right next step, while Maya stays in control of what goes out in her name?

Solution and proof

Reply-aware sequencing: every reply gets a next step, and the SDR decides how much the AI decides.

Shape

Measure

Prototype

Reply-aware sequencing, in five moves



Inv1

- Intent detection on every reply, across six intents
- Branch library: not now resumes on their timeline, referral drafts a warm intro, out of office pauses
- Approval queue with reason and confidence
- Autonomy dial per intent
- Hard exit on opt out and clear negatives

Next

- Pooled message testing, judged on meetings
- Branching on pre-reply signals
- Thread summary for the AE

Not now

- A contact database or enrichment
- Own sending and warmup
- A fully autonomous AI SDR

One revenue number, guarded on four sides

NORTH STAR

**Meetings booked per
100 prospects who reply**

+20%

versus control, over an 8-week pilot with 5 to 8 SDR teams. Counts only meetings the AE accepts.

Trust

Suggestions approved without edits

70%+ by week 6

Intents switched to auto

2 of 6 by week 8

Speed

Median reply to next touch

Under 4 working hours

Manual reply handling

Down 50%

Quality

Intent accuracy, weekly audit of 100

90%+

Guardrails

Spam complaint rate

Below 0.1%


Unsubscribe rate

No higher than control


Sends after an opt out

Zero


Four screens, one reply, end to end




SCREEN 1
Reply inbox
Every reply tagged with intent and confidence, sorted by what needs her first



SCREEN 2
Branch card
The suggested step, the draft and the reason, with approve, edit, skip



SCREEN 3
Autonomy settings
A dial per intent, from review everything to auto



SCREEN 4
Pilot view
Meetings per 100 repliers against control, plus guardrails

THE PRD COVERS

Problem and P0 evidence / persona and JTBD / scope / requirements per intent / edge cases / AI behavior and thresholds / human-in-the-loop rules / integrations with HubSpot, Apollo, Gmail, Outlook / metrics / rollout plan

RISKS DESIGNED AROUND

- **Wrong intent:** every intent starts in review, auto is opt-in
- **Shallow APIs:** start with HubSpot and Apollo
- **Thin data:** pool 8 weeks, confirm on accepted meetings

SOURCES

Secondary research for acts two and three

1. The Bridge Group, 2025 SDR Models and Metrics Report. bridgegroupinc.com
2. Salesforce, State of Sales, sales AI statistics. salesforce.com
3. Gartner, 61% of B2B buyers prefer a rep-free buying experience, Jun 2025. gartner.com/en/newsroom
4. Gartner, AI agents will outnumber sellers by 10x by 2028, Nov 2025. gartner.com/en/newsroom
5. Unify, How to A/B test outbound with small sample sizes. unifygtm.com
6. Instantly, Cold email reply rate benchmarks, citing Belkins and Backlinko. instantly.ai/blog
7. Deepline, GTM stack Q&A on AI SDR results. deepline.com/gtm-stack
8. RevBoss, Eric Boggs, On AI SDR-ing. email.revboss.com
9. Broadn, Why AI SDRs are doomed to fail. broadn.io/blogs
10. Close, Requirements for bulk email senders by Google and Yahoo. help.close.com

G2 reviews used in act one were intentionally excluded here. Vendor-published figures (Instantly, Unify) are directional benchmarks only. The persona, journey and scoring are my synthesis of these sources, and pilot targets are goals I set for this concept.

THANKS FOR READING

Tell me where I'm wrong.



I'm studying how real teams find, win and keep customers. If you sell, run RevOps or build revenue tools, I'd love to hear how you handle the reply.

[linkedin.com/in/thakursmrity](https://www.linkedin.com/in/thakursmrity)

sequorai.com