

BENCHMARK REPORT • 2026 EDITION

The Multichannel Outreach Report 2026

Email

+

LinkedIn

+

Voice

What the numbers say about signals, send windows and suppression, and how to build a sequence that earns a reply when email alone no longer does.

SIXTEEN PUBLISHED DATASETS, COMPILED AND ATTRIBUTED

01 BASELINE | 02 LENGTH | 03 CHANNELS | 04 SIGNALS | 05 DECAY |
06 CADENCE | 07 ARCHITECTURE | 08 SELF AUDIT

3.43%AVERAGE COLD EMAIL
REPLY RATE, 2026**5 to 18%**REPLY RATE ON SIGNAL LED
OUTREACH**2 to 3x**

LIFT FROM COORDINATED CHANNELS

47hrsAVERAGE TIME TO ACT ON
A SIGNAL**Ishan Vats**

Founder, IV Consulting · Clay and outbound systems

COMPILED
SEPTEMBER 2026

HOW TO READ THIS

Sixteen published datasets, one argument: **the channel is not the problem.**

Cold email reply rates have roughly halved in two years while every sending platform got better. That is not a copy problem and it is not a tooling problem. It is what happens when the same list gets contacted on one channel by everybody at once.

What is in here

- 01 The baseline.** Where reply rates actually sit, and the deliverability gates you now have to clear to be measured at all.
- 02 Message length.** The word band that outperforms, with the reply curve, and how deep the sequence should run.
- 03 Channel by channel.** Email, LinkedIn and phone side by side on the same scale.
- 04 Signals.** The lift from trigger led outreach, and which triggers carry it.
- 05 Decay.** How fast a signal stops being worth acting on, by trigger type.
- 06 The build.** The cadence, the suppression rules, a reference architecture and a twelve box self audit.

On the numbers

Every figure here is somebody else's published data, attributed on the page it appears on, with the full list on page 12. We have not restated anybody's research as our own and we have not adjusted a number to make a point.

Methodologies differ, and a reply is counted differently across all of them. Where sources disagree we give the range rather than pick the flattering end.

What is ours

Pages 8 to 10 are how we build: the cadence, the three rules that hold it together, the suppression model and the reference architecture. They are marked as ours wherever they appear.

Treat them as a starting point to instrument against your own data, not a benchmark to hit.

If you read one page

Page 11. Eight lines, in the order we would act on them, each one tied back to the figure on the page it came from.

Page 10 is the twelve box audit if you would rather start by finding out where you stand.

The teams still growing outbound in 2026 are not sending more. They are sending for a reason, to fewer people, across more than one channel, and they stop the moment somebody answers.

01 • THE BASELINE

Reply rates have **halved** in two years.

3.43%AVERAGE REPLY RATE,
2026**5.1%**THE SAME FIGURE IN
2024**10%+**WHERE ELITE SENDERS
RUN**0.3%**

SPAM COMPLAINT CEILING

INSTANTLY 2026 BENCHMARK REPORT • LEMLIST • AMPELMARKET 2026 COLD EMAIL BENCHMARKS

The decline is enforcement, not fatigue

The drop tracks three things at once: inbox saturation, tighter Gmail and Outlook spam enforcement, and a flood of low effort AI generated outreach that trained both filters and recipients to pattern match it. Sending more of the same is the one response that reliably makes it worse. lemlist's own data across 350M emails and 4.5M LinkedIn messages finds higher volume correlating with lower reply rates.

Which means the gates on the right are not deliverability hygiene any more. They are the price of having your campaign measured at all. A campaign that misses them is not underperforming, it is unmeasured.

Read the bands rather than the average. 1 to 5% is the standard B2B band, 5 to 10% is a good campaign, and above 10% is where tight targeting puts you. Nothing in this report will move a bad list into the top band. Everything in it is aimed at earning the difference between the first band and the third.

GATE	ENFORCED	WHERE TO RUN
Spam complaints	under 0.3%	under 0.1%
Bounce rate	under 2%	under 1%
SPF, DKIM, DMARC	required	aligned
One click unsubscribe	RFC 8058	2 day window
Inbox placement	not published	80% or better
Volume	not published	capped per inbox

What this page costs you if you skip it

Everything after here assumes the gates are clear. Signal detection, scoring and multichannel routing all improve a number that a spam folder makes zero. Fix the infrastructure first, because it is the cheapest work in the report and the only work that nothing else can compensate for.

SOURCES: INSTANTLY, LEMLIST, AMPELMARKET, WARMFORGE, GOOGLE AND YAHOO BULK SENDER REQUIREMENTS

02 • MESSAGE LENGTH AND DEPTH

The reply curve peaks between 50 and 125 words.

Across more than 4 million emails analysed by Boomerang, lemlist, Instantly and Gong, reply rate peaks in the 50 to 125 word band and falls away steadily from there. A separate study of 3 million cold emails puts the same band at 2.4 times the reply rate of anything over 200 words.

REPLY RATE BY WORD COUNT



BOOMERANG, LEMLIST, INSTANTLY AND GONG, 4M+ EMAILS • OVERLOOP 2026

Why short wins

A cold email is read on a phone, in a list, in about four seconds. Anything past the first screen is a decision to keep reading that you have not earned yet.

What short is not

Short is not vague. The band that wins carries a specific reason for writing, to this person, this week. Cutting the reason to hit a word count removes the only thing doing work.

The practical shape

One line of why now, one line of what you do for companies in that position, one question answerable in a word. That lands between 60 and 90 words without trying.

4 to 7

EMAILS IN THE BAND THAT MAXIMISES REPLY RATE

58%

OF REPLIES ARRIVE ON STEP ONE

42%

ARRIVE ON STEPS TWO TO FOUR

14 to 21

DAYS IS THE WHOLE USEFUL RUNWAY

The follow up carries more than you think

58% of replies to a cold campaign arrive on step one. The other 42% arrive on steps two to four, which means a two step sequence gives away four replies in ten before it starts. Three to five emails over 14 to 21 days captures 78 to 90% of the replies you were ever going to get, without the complaint rate that longer sequences generate.

INSTANTLY 2026 BENCHMARK REPORT

03 • CHANNEL BY CHANNEL

Three channels, three different jobs, **one account.**

Put the published benchmarks on the same page and the case for coordination makes itself. No single channel is strong enough to carry a number on its own, and each one fails in a different direction.



CHANNEL 01

Email

Average reply

3.43%

Strong

above 5%

Elite

8 to 12%

Open, raw list

15 to 25%

Best length

50 to 125 w

Scales furthest, degrades fastest. Its failure mode is silent: you keep sending and the filter keeps quietly deciding.



CHANNEL 02

LinkedIn

Acceptance, cold

21 to 28%

Strong

above 35%

Personalised

about 45%

Generic

about 15%

Fix targeting below

20%

Highest intent per touch, lowest ceiling on volume. Acceptance is a targeting readout before it is a copy readout.



CHANNEL 03

Phone

Connect, generic

5 to 12%

Verified mobile

18 to 22%

Dials per meeting

25 to 35

Top performers

12 to 18

Dial to meeting

2 to 3%

Slowest and most expensive per touch, and the only one that gets you an answer today. Data quality is the whole lever.

The number that makes the argument

Coordinated email, LinkedIn and phone against the same account produces **2 to 3 times** the replies of any one channel run alone. Not because three messages beat one, but because the second and third touch arrive carrying context the first one established.

LANDBASE, VIA LEMLIST MULTICHANNEL OUTBOUND ANALYSIS

The corollary nobody prices in

Three uncoordinated channels hitting one person in the same week is not multichannel. It is three people ignoring each other inside your own company, and the prospect experiences it as one company that does not know who has already contacted them.

Where the lift actually comes from

The second touch works because the name is now familiar and the reason is already on the record. That only holds if the channels run in a fixed order against one account, with one shared view of what has been sent. Page 9 is the rule set that produces it.

Adding a channel multiplies your reply rate. Adding a channel without shared state multiplies your complaint rate.

04 • SIGNALS

The same message, sent for a reason, replies **five times better.**

This is the single largest effect in the published data, and the only one that does not require you to send more. Generic cold email sits at 1 to 3%. The same outreach attached to a trigger runs at 5 to 18%.

TRIGGER	REPORTED REPLY OR RESPONSE	AGAINST A BASELINE OF	WINDOW
Job change, decision maker	18% or higher	1 to 5%	First 90 days
Leadership change at the account	14%	1.2% cold call	First 100 days
Former champion in a new role	6.3x conversion	cold prospect	On detection
Funding round announced	5 to 18% band	1 to 3%	2 to 4 weeks
Tech stack change	5 to 18% band	1 to 3%	On detection
Hiring for the function you sell to	5 to 18% band	1 to 3%	While the role is open

FIRSTSALES 90 DAY WINDOW STUDY • UNIFY GTM LINKEDIN SIGNAL RANKING • AUTOBOUND SIGNAL BASED SELLING 2026 • LEAD-SCORER TRIGGER STACK ANALYSIS

70%

OF A NEW EXECUTIVE'S BUDGET IS SPENT IN THEIR FIRST 100 DAYS

6.3x

CONVERSION ON A FORMER CHAMPION AGAINST A COLD PROSPECT

18%+

REPLY RATE INSIDE THE 90 DAY JOB CHANGE WINDOW

A signal is a reason, not a merge field

The lift does not come from writing "congratulations on the new role". It comes from the fact that the person genuinely has a new problem this month, and you are the only message in their inbox that is about it. Drop the trigger into a template and you keep the tokens and lose the effect.

Which is why signal quality beats signal quantity. Six triggers you can write a genuinely specific first line from will outperform sixty you can only reference.

HOW WE BUILD IT

Score, then route

Every account carries one composite score built from trigger type, how old the trigger is, fit, and whether we have a verified mobile. The score decides the channel order, not a rep's preference: high score with a verified mobile routes to phone first, high score without one routes to LinkedIn, everything else starts on email.


If you change one thing after reading this report, change what causes a send. **Not the copy, not the tool. The reason.**

05 • DECAY

A signal has a half life, and most teams miss it.

Response rate to signal triggered outreach falls by roughly 40 to 60% after the first 48 hours. The average company takes 47 hours to respond to a lead at all.

RESPONSE RATE BY TIME SINCE THE SIGNAL FIRED



BOUNCE WATCH FIRST 24 HOURS ANALYSIS • OPTIFAI BENCHMARK OF 939 B2B COMPANIES, 2025 TO 2026

The trap in the numbers

47 hours is the average, and the decay starts at 24. So the median signal is acted on after its best window has already closed, by a team that will then conclude signals do not work and go back to volume. The data does not say signals are weak. It says slow signals are weak.

What fast has to mean operationally

Under 24 hours from detection to first touch, which rules out a human deciding. It has to be a scheduled job: the signal lands, enrichment runs, the score is written, the account is routed, and the first touch queues without anybody opening a tab. A weekly list pull cannot hit this window no matter how good the list is.

SLOW DECAY

Job change

A 90 day window. Fast still wins, but you have weeks rather than hours to work with.

MEDIUM DECAY

Funding

Act inside 2 to 4 weeks, while the spending plan is still being written rather than executed.

FAST DECAY

Site or ad behaviour

Hours, not days. This is the one that genuinely dies overnight, and the one worth automating first.

One number to instrument this week

Median hours from signal detected to first touch sent. Most teams have never measured it and cannot say whether theirs is 6 or 96. It is a single timestamp subtraction, it needs no new tooling, and it will tell you more about your outbound than any reply rate on this page.


WINDOWS PER FIRSTSALES, BOUNCE WATCH AND AUTOBOUND. TREAT THEM AS STARTING SLAS, NOT TARGETS.

06 • THE BUILD, PART ONE

The cadence, and the three rules that hold it together.

Everything so far was somebody else's data. From here it is how we assemble it, which is ours and is marked as ours. Fourteen days, one account, six touches, in a fixed order.

OUR REFERENCE CADENCE, ONE ACCOUNT, 14 DAYS



Three rules that do the work

- 01 **One channel per account per day.** Never two touches on the same person inside 24 hours, on any channel, from anybody in the company.
- 02 **A response suppresses everywhere.** A reply, an accept, an answered call or a form fill stops every other queued touch for that account immediately, including a colleague's.
- 03 **Every step carries new information.** If a step exists only to move the thread to the top of the inbox, delete it. That step is where complaint rate comes from.

Why voice sits at day 4

Late enough that the name is familiar from two prior touches, early enough that the signal has not decayed. Connect rate on a verified mobile runs 18 to 22% against 5 to 12% on generic data, so the step only exists where the number is verified. Everywhere else it is skipped, not guessed.

The failure this prevents

The common multichannel build is three tools running three lists with no shared state. The prospect gets five touches in a week, two of them after they already replied. That is not a 2 to 3x lift, it is a complaint and a burned domain.

Day 14 is an exit, not a last try

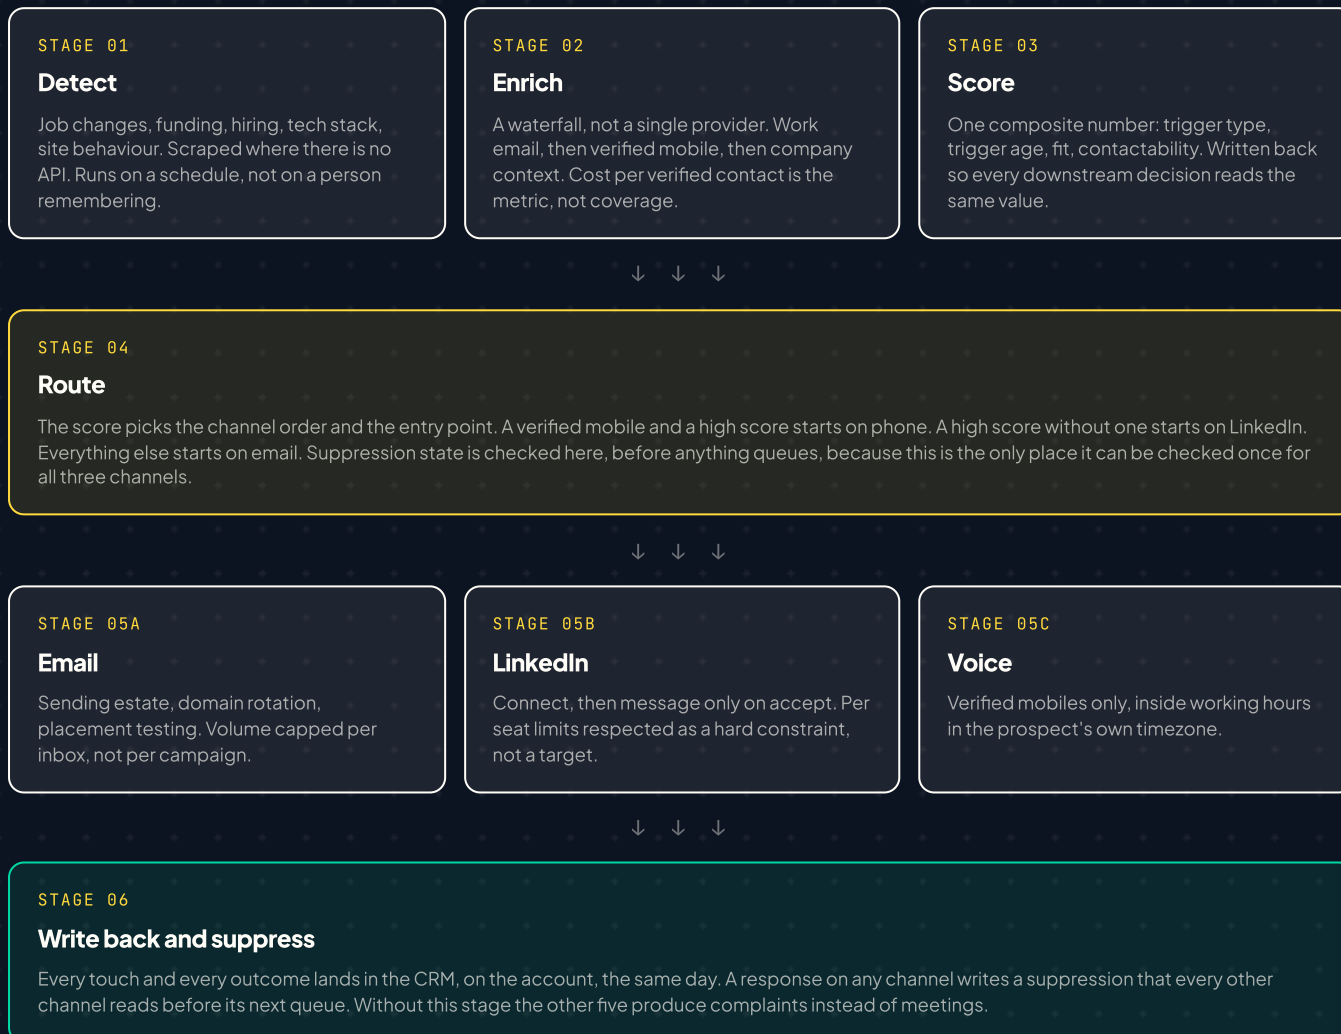
One line, no guilt, no fourth bump, and the account goes back to the pool to wait for its next signal. The sequence ends because the reason expired, not because you ran out of templates.

Six touches in fourteen days is not a volume decision. It is what fits inside one signal before the reason you were writing stops being true.

07 · THE BUILD, PART TWO

The reference architecture.

Six stages. Most teams have the last two, which is exactly why their signals arrive at 47 hours and their channels do not know about each other.



The one thing to take from this page

Suppression is not a feature of your sending tool. It is shared state that has to live outside all three channels, because no channel can see the other two. That is the whole difference between multichannel and three channels.

08 • SELF AUDIT

Twelve boxes. Tick what is true today.

Not a maturity model. Each line is something either running or not running in your estate right now, and each one maps to a number earlier in this report.

- | | |
|--|--|
| ☐ Authentication. SPF, DKIM and DMARC are aligned on every sending domain. | ☐ Complaint rate. You know this week's number and it is under 0.1%. |
| ☐ Placement. You test inbox placement, not just opens and bounces. | ☐ Length. Your first email is between 50 and 125 words. |
| ☐ Depth. Your sequence runs four to seven steps, not two. | ☐ New information. No step exists only to bump the thread. |
| ☐ Triggers. At least three signals are detected on a schedule, not pulled by hand. | ☐ Speed. Detection to first touch is under 24 hours, and you have measured it. |
| ☐ Scoring. One composite score exists and every channel reads it. | ☐ Routing. The score picks the channel, not the rep's preference. |
| ☐ Suppression. A reply on one channel stops the queue on the other two, same day. | ☐ Memory. Every touch and outcome is written back to the CRM automatically. |

0 to 4 ticked

You are running volume. The fastest gain is not more sending, it is the deliverability gates and the length band. Both are free and both are this week's work.

5 to 8 ticked

You have a working channel and no system. The gap is detection and scoring, which is where the 5 to 18% band lives and where the compounding starts.

9 to 12 ticked

You have the architecture. The remaining work is measurement: cost per verified contact, and reply rate by trigger type rather than by campaign.

The four numbers most teams cannot produce

Complaint rate this week. Median hours from detection to first touch. Reply rate split by trigger type rather than by campaign. Cost per verified contact. Everything else in the audit is a build decision you can argue about. These four are a readout, and you either have them or you are guessing.

09 • THE SHORT VERSION

Eight lines, if you only remember **eight lines**.

- 01 **Email alone is a declining asset.** 5.1% in 2024, 3.43% in 2026. Sending more of it is the response that reliably makes it worse.
- 02 **The gates are now the entry fee.** Complaints under 0.3%, bounces under 2%, authentication aligned, one click unsubscribe. Miss these and nothing else you do gets measured.
- 03 **Write between 50 and 125 words.** Reply rate peaks there at 8.2% and falls to 3.9% past 200 words.
- 04 **Go four to seven steps.** 42% of replies arrive after step one. A two step sequence gives them away before it starts.
- 05 **Attach every send to a reason.** Signal led outreach runs 5 to 18% against 1 to 3% generic. This is the largest effect in the data and the only one that does not cost volume.
- 06 **Act inside 24 hours.** Response rate falls 40 to 60% after 48. The average company takes 47, which is why most teams conclude signals do not work.
- 07 **Coordinate, do not just add.** Coordinated channels return 2 to 3 times a single channel. Uncoordinated ones return complaints.
- 08 **Suppression lives outside the channels.** No sending tool can see the other two, so shared state is the whole difference between multichannel and three channels.

What we would do first, in this order

Fix the gates. Cut the first email to the word band. Add steps three and four. Pick the three triggers you can genuinely write a specific first line from and detect them on a schedule. Put a score on the account. Then, and only then, add the second and third channel, with suppression built before the first send goes out.

SOURCES AND METHOD

Where every number came from.

This is a compilation, not primary research. We did not run these studies and we do not present them as ours. Figures are reproduced as published, ranges are given where sources disagree, and nothing has been adjusted.

Instantly, Cold Email Benchmark Report 2026

instantly.ai/cold-email-benchmark-report-2026: platform average reply rate, the 2024 to 2026 decline, sequence length bands, reply distribution by step.

lemlist, cold email response rates and multichannel outbound analysis

lemlist.com/blog: 350M emails and 4.5M LinkedIn messages, volume against reply rate, the Landbase 2 to 3x coordinated channel figure, reported multichannel campaign results.

Boomerang, lemlist, Instantly and Gong, combined length study via overloop.com/blog/whats-the-best-email-length-for-sales-outreach: reply rate by word count across 4M+ emails.

Amplemarket, 2026 cold email benchmarks

amplemarket.com/blog/cold-email-benchmarks: bounce, open, reply and spam rate bands on managed sending infrastructure.

Warmforge, email deliverability statistics 2026

warmforge.ai/blog/email-deliverability-statistics: complaint and bounce thresholds, inbox placement targets.

Google and Yahoo bulk sender requirements

Authentication, RFC 8058 one click unsubscribe, the 0.3% complaint ceiling and the two day unsubscribe window.

Expandi, State of LinkedIn Outreach 2026

expandi.io/state-of-linkedin-outreach-h2-2026: acceptance rate bands across 13.2M data points, the below 20% targeting warning. HeyReach campaign dataset referenced for the cold B2B skew.

Overloop, LinkedIn outreach benchmarks 2026

overloop.com/blog/linkedin-outreach-benchmarks: personalised against generic acceptance rate.

Skipcall, cold call connect rate benchmarks 2026

skipcall.io/en/blog/cold-call-connect-rate-benchmarks: connect rate on generic against verified mobile data.

Salescadia and Martal, cold calling benchmarks 2026

Dials per booked meeting, dial to meeting conversion, the effect of data quality on connect rate.

Firstsales, the 90 day job change window

firstsales.io/blog/job-change-trigger-email: reply rate inside the window, new executive budget behaviour in the first 100 days.

Unify GTM, LinkedIn signals that predict outbound conversion

unifygtm.com/explore/linkedin-signals-predict-outbound-conversion: trigger ranking, former champion conversion multiple.

Autobound, signal based selling 2026

autobound.ai/blog/signal-based-selling-complete-guide: signal led reply bands, funding and tech stack windows, leadership change response rate.

Lead-scorer, the 2026 trigger stack

lead-scorer.com/blog/buying-signals-b2b-2026: the 3.4% to 18% framing.

Bounce Watch, the first 24 hours after a signal

bouncewatch.com/blog/signal-intelligence/first-24-hours-signal-speed-outreach: response rate by hours since the signal, the 40 to 60% decay after 48 hours.

Optifai, speed to lead benchmark

939 B2B companies, 2025 to 2026, via the digitalapplied speed to lead analysis: the 47 hour average response time.

How to use a compiled benchmark

Platform datasets skew toward cold lists sent at volume, agency datasets skew warmer, and a reply is defined differently across all of them. Use these as a direction of travel and a sanity check on your own instrumentation. The only benchmark that should change a decision is your own, measured the same way twice.

WHAT TO DO WITH IT

Reading the report is not the same as building the system.

The Outbound Build is a 6 week program where you build exactly this architecture, on your own account, across twelve live sessions. Detection, waterfall enrichment, composite scoring, three channels, shared suppression, CRM writeback. You leave with it running, not with notes about it.

6

WEEKS

12

LIVE SESSIONS

10

TOOLS IN THE ESTATE

1

LIVE ACCOUNT, YOURS

Application only, and deliberately selective. Four gates: a real account to build on, the authority to run outbound on it, the volume to learn from, and the time to attend live.

Applications for Cohort 2 are openivconsulting.in/gtm-engineering-programTHE OUTBOUND BUILD
TAUGHT BY ISHAN VATS