

The State of Cold Email 2026

Smartlead • 2026 Edition



Contents

A note from Smartlead	01
How we built this report	02
Key findings	03
Where you stand	04
The behaviors behind the gap	06
Getting into the inbox	09
What earns the reply	12
Signals worth knowing	14
About Smartlead	16
Appendix	17



A note from Smartlead

Something has quietly changed in cold email, and it showed up in our data before it showed up anywhere else.

The best senders earn a reply for roughly every 38 people they contact, while the typical sender needs around 135, and the weakest can go past 600. They all send into the same inboxes over the same six months, yet the difference in results is enormous.

Most reports on cold email tell you what the average sender does, but an average hides the one thing that matters most, which is the huge spread between the best senders and the worst. So this report is built to answer a different question. Not what is normal, but where do you stand.

Every chart shows the full range, from the strongest senders down to the weakest, so you can place yourself on it and see how far ahead the leaders are. Because more than 850 million emails were sent through Smartlead in the first half of 2026, we can measure that range directly instead of guessing at it from a survey.

We report what the data shows, including where it disagrees with common advice and where we simply cannot measure something yet. This edition covers cold email, from the moment you send to the moment someone replies.



How we built it

This report is built from combined, anonymous sending data on the Smartlead platform, covering real sends, deliveries, and replies. We remove test and internal accounts, and no individual business or recipient can be identified in any finding.



Evidence base

More than 850 million emails sent through Smartlead



Time period

6 months Jan to June 2026



Trend view

Month by month across the first half of 2026



Data source

Smartlead platform only, no outside research

How to read the numbers

Most of the reply charts use one simple idea: how many people you contact to earn a single reply, where a lower number is better.

The behavior charts line up every sender from weakest to strongest and mark six points along that line. P10 is near the bottom, the point only 10% of senders fall below. P25 is the lower quarter, P50 is the middle, the typical sender, P75 is the upper quarter, P90 is near the top, the point that only the best 10% beat, and P95 is the very top. To use a chart, find your own value, see which point you land near, and see how far the leaders are from there.



Three things the data made clear

This report is built from combined, anonymous sending data on the Smartlead platform, covering real sends, deliveries, and replies. We remove test and internal accounts, and no individual business or recipient can be identified in any finding.

~38

Contacts is all it takes the top 10% to earn a reply.

The best senders earn a reply for roughly every 38 people they contact, and the very best need only 26. The typical sender needs 135, while the weakest go past 600. That is about a sixteen times difference in results for reaching the same kinds of inboxes.

3.5x

Separates the typical sender from the top 10%, and it held all half year.

Through the first half of 2026 the typical sender stayed around 135 contacts per reply while the top 10 percent stayed near 38, a gap of about three and a half times that barely moved from month to month. This gap is built into how people send, not a seasonal blip.

2%

Bounce rate is the line most senders should never cross.

A bounce is an email that fails to deliver, and too many of them get your future emails filtered out. The typical sender bounces about 1.54 percent of the time, which stays under the danger line, but the weakest quarter of senders sit above 3 percent. Clean lists and steady sending are what separate the two groups.






Where you stand

Something has quietly changed in cold email, and it showed up in our data before it showed up anywhere else.

Across the platform, the effort behind a single reply varies enormously. The most efficient senders earn one for every few dozen people they contact, while the least efficient work through hundreds. The two charts below place you in that range, first as it stands in the first half of 2026, and then month by month across the half year.

CHART 1 The top 10% earn a reply every 38 contacts, while the bottom 10% need more than 600.

Contacts needed per reply, by sender group · lower is better



Find your place on this chart.


The very best senders earn a reply for every 26 people they contact, and the top 10 percent for every 38. The typical sender needs 135, and below the middle the effort climbs steeply, with the lower quarter needing about 270 and the bottom 10 percent more than 600. The distance between the typical sender and the best is the whole opportunity.

Contacts per reply is simply one divided by the reply rate.

CHART 2

Through the first half of 2026, the typical sender's effort held steady near 135 contacts per reply.

Contacts per reply, month by month · typical sender against the top 10 percent · lower is better



The picture was steady rather than improving.

From month to month the typical sender moved only between about 123 and 147 contacts per reply, drifting slightly harder as the year went on, while the top 10 percent stayed near 27 to 33. There was no seasonal easing in the first half, and the gap that opened in earlier years simply continued.





The behaviors behind the gap


List quality, list size, how much you send, how many mailboxes you use, and how much you personalize. These are the choices that separate the top of the range from the bottom.

Chart 1 shows how far apart senders are, and these five charts show what tends to move with that gap. Each one shows the full range across senders in the first half of 2026, so you can see not only the typical value but where the best and worst actually sit.

CHART 3

List quality shows up first in bounce rate, and a quarter of senders are in the danger zone.

Bounce rate across senders · lower is better



The typical sender bounces about 1.54 percent of the time, which is roughly one failed email in every 65.

The safe limit before email providers begin blocking your future messages is around 2 percent. The typical sender stays under it, but the upper quarter of senders, sitting at 3.21 percent and above, are over the line. Bounce rate is usually the first place list quality shows up, with clean and verified lists near the bottom and cheap, bulk bought lists near the top.


For comparison, Gmail averages about 2.95 percent bounce and Outlook about 4.34 percent since Microsoft's May 2025 changes.



CHART 4

The typical campaign contacts 499 people, while the biggest go past 5,700.

People contacted per campaign · list size




The typical campaign contacts 499 people.

In the data, larger lists tend to earn replies at or below the typical rate, while smaller and more focused campaigns tend to do better. Bigger is not better, because a small list aimed at the right kind of buyer usually earns more than a large, loosely chosen one.

CHART 5

The heaviest senders push each mailbox close to the safe limit.

Emails sent per mailbox per day · sending pace



The typical sender sends about 14 emails from each mailbox per day.

The widely used safe limit is 50. The heaviest 10 percent sit at 42, just under it, and the very heaviest reach about 70, well over, which probably feeds the high bounce rates in Chart 3. The most efficient senders do not push each mailbox the hardest. They keep the pace per mailbox moderate and add more mailboxes instead.


Worked out for each sender as emails sent divided by mailboxes used divided by days they sent.



CHART 6

The typical account runs 20 mailboxes, while agencies run hundreds.

Mailboxes connected per account · size of the sending setup




The typical account runs 20 mailboxes.

Using more mailboxes spreads the sending around, so no single one builds up enough bad marks to get blocked. Size on its own does not decide results. The best senders for replies do run more mailboxes than average, but they combine that with the lower pace per mailbox from Chart 5, rather than pushing each one harder.

CHART 7

Personalization splits the market. Most senders do it on every email, and a quarter barely do it at all.

Emails sent per mailbox per day · sending pace



Senders gather at the two extremes.

The bottom 10 percent personalize nothing, the lower quarter about 14 percent, and from the upper quarter onward it reaches 100 percent. The typical figure of 75 percent reflects a split crowd rather than a middle habit, because the weakest senders for replies tend to personalize nothing while the strongest personalize every email.

Based on whether a sender's first emails add the recipient's own details.





Getting into the inbox


Proving your emails are genuine, choosing where you send from, and warming up your mailboxes. These are the basics that decide whether anyone ever sees your email.

New email rules from Google and Yahoo in February 2024, followed by Microsoft in May 2025, raised the bar for reaching the inbox. We cannot see exactly where every email landed, but the data is clear on the things you control, which are how you prove your emails are genuine, which provider you send from, and how you warm up your mailboxes.

CHART 8

Proving your emails are genuine saves you roughly 1 in 9 of your contacts.

Contacts per reply, by whether the sending setup is verified · lower is better



Set up the three standard checks that prove your email is really from you, known as SPF, DKIM, and DMARC.

Campaigns that pass all three earn a reply every 108 contacts, while those that do not need 120, which is about 11 percent more work for the same result. It is a one time setup and the cheapest improvement available.


Measured per campaign where the sending setup is known, which is about half of all volume. We use replies as the standard, because we cannot directly see where each email landed.



CHART 9

Gmail earns replies with the least effort, while custom servers need the most.

Contacts per reply, by the provider you send from · lower is better



The provider you send from changes the result.


Gmail earns a reply every 90 contacts, Outlook every 120, and custom mail servers every 179. Outlook also has the highest bounce load, at about one failed email in 23 compared with one in 34 on Gmail, which is a sign of tougher filtering since Microsoft's May 2025 changes.

Based on campaigns that use a single provider, which is about 40 percent of volume. Large campaigns that mix providers are left out.

CHART 10

During warmup, every major provider reaches the inbox about 99 percent of the time.

Warmup inbox placement rate · higher is better



The warmup system is healthy.

Outlook edges ahead at 99.04 percent, with custom servers and Gmail at 98.83 percent. These are close because this is warmup mail, the friendly emails that Smartlead mailboxes send each other to build trust, and not real cold campaigns. Treat it as a health check rather than a target to aim for.


The chart starts at 98 percent so the small differences are visible.



CHART 11

Senders are warming up faster, with far less idle time.

Days from connecting a mailbox to the first campaign



The typical wait held steady at 8 to 9 days, while the average dropped sharply.

The time from connecting a mailbox to launching the first campaign stayed near 8 to 9 days, moving from 9.4 to 8.0. The average fell from 28.6 to 19.8 days, which means far fewer accounts now sit unused for months before launching.

These are yearly groups by signup date, because giving each account a full year to launch is the fairer measure. A six month group would unfairly leave out accounts that signed up late and simply have not launched yet. Launch means the first campaign was created.






What earns the reply

How many follow ups you send, how long your sequence runs, and how long your emails are. These are the choices about the message itself, settled here by the data.

CHART 12 Your first email earns a reply with about a third of the effort of your seventh.

Contacts per reply, by which email in the sequence it is · lower is better



The first email does most of the work.


The first email earns a reply every 85 contacts, and by the seventh it takes 303, which is about three and a half times the effort for the same reply. People rarely unsubscribe from later emails. They simply stop reading, so the cost stays hidden until the results quietly slip.

The unsubscribe rate stays flat at 0.02 to 0.05 percent across every email in the sequence.

CHART 13

The typical campaign runs 3 emails, and the bottom 10% send only one.

Number of emails in a campaign's sequence




The typical campaign runs 3 emails, while the bottom 10 percent send just one.

Read alongside Chart 12, those single email campaigns give up all the replies a second email would have earned. The gains fade quickly, though, because past the fifth email each one earns a reply only about every 270 contacts or worse, so adding more is not automatically better.

CHART 14

Short emails earn replies with the least effort.

Contacts per reply, by the length of the first email · lower is better



Short emails win.

First emails under 50 words earn a reply every 68 contacts, the best of any length, while emails in the padded 100 to 149 word range need 97. The typical sender writes about 98 words, which sits right at the edge of that weaker range, so most senders would do better by cutting.

Word count is taken from the first email, with formatting and inserted details removed.





Signals worth knowing


Smaller patterns from the platform that are worth keeping in mind.

One note on reading these, and the patterns in Section 2. They are patterns in the data, not controlled tests. A strong sender often does several of these things at once, so treat them as links worth testing rather than proof that one thing causes another.

CHART 15

Sending early beats sending into the crowd.

Contacts per reply by the hour you send (in UTC), with how busy each hour is · lower is better



The best window is 6 to 10 in the morning UTC


with a reply every 75 contacts at the 7am peak. The busy stretch from 1 to 5 in the afternoon UTC, when most senders send, needs about 120. Sending at the same time as everyone else makes every reply cost more. Times are in UTC, not the reader's own local time.



CHART 16

The best day to send is mostly a myth, and even Saturday is close.

Contacts per reply by day of the week, with how busy each day is · lower is better




Weekdays barely differ

all needing about 110 to 115 contacts per reply, with Wednesday slightly ahead at 110. The weakest day is Saturday at about 122, which is a smaller gap than usual. Sunday holds at about 111 and matches the weekdays, even though far fewer emails go out then, which makes it a quiet window most senders skip.

CHART 17

If a reply has not come within a day, it almost never will.

Share of replies that arrive within each window of time



Replies arrive fast

with 64 percent coming in the first hour, about 89 percent within a day, and 94 percent within three days. The typical reply lands in about 10 minutes. Follow ups sent days later are mostly reaching people who have already decided not to respond.

Based on 7.9 million replies. Measured as the reply time minus the send time.



About Smartlead

We produced evidence based analysis of cold email, built from combined, anonymous activity, meaning real campaign sends, replies, and delivery results, rather than surveys. Smartlead is a cold outreach and sending platform with AI-native outbound suite across Email and Calling used by more than 100,000 businesses around the world.



A. Definitions and method

Contacts per reply

This is how many people you contact to earn one reply, and a lower number is better. It is simply the reply rate turned around, so 135 contacts per reply is the same as a 0.74 percent reply rate. It scales in a straight line, so to land 10 replies at the typical rate you would plan for about 1,350 contacts.

Reply

A genuine reply from a person, counted against delivered emails, leaving out automatic out of office messages.

The percentile points

If you line up every sender from weakest to strongest, P10 is the point that only 10 percent fall below, so it sits near the bottom. P25 is the lower quarter. P50 is the middle, the typical sender. P75 is the upper quarter. P90 is near the top, beaten by only the best 10 percent, and P95 is the very top. Only senders with at least 500 emails sent in the period are included.

The behavior charts

The ranges in Section 2, covering bounce, list size, sending pace, mailboxes, and personalization, along with the sequence length chart, come from active senders and campaigns during January to June 2026.

Open rate

We deliberately do not use it. Apple Mail loads email images automatically and triggers open tracking before anyone reads the message, so reported open rates of 60 to 70 percent are largely automatic. Only a reply shows genuine interest.

Campaign level figures

Two charts use campaign level data because the detail is not available for each single send. These are whether the sending setup is verified (Chart 8, about half of volume) and which provider was used (Chart 9, single provider campaigns, about 40 percent of volume). Treat those two as pointing in the right direction rather than exact.

Two charts on purpose

Two charts are not limited to the six month window. Warmup inbox placement (Chart 10) uses warmup figures gathered across the whole platform over time, because dated warmup records are not available, and the number barely moves at around 99 percent anyway. Days to first campaign (Chart 11) uses yearly signup groups on purpose, because a six month group would unfairly leave out accounts that signed up late and have not launched yet, which would drag the typical figure down, so the full year view is the fairer measure.



B. The numbers behind the charts

Reply efficiency by sender group (Chart 1), contacts per reply

Group	Top 5%	Top 10%	Top 25%	Typical	Bottom 25%	Bottom 10%
Contacts per reply	26	38	70	135	270	625

Behavior ranges (Charts 3 to 7, and 13), by point

Measure	P10	P25	P50	P75	P90	P95
Bounce rate	0.28%	0.73%	1.54%	3.21%	6.71%	10.85%
People per campaign	25	87	499	1,961	5,775	9,871
Emails per mailbox per day	4.4	8.1	13.9	23.5	42.1	70.2
Mailboxes per account	3	7	20	60	214	535
Personalization usage	0%	14.3%	75%	100%	100%	100%
Emails in the sequence	1	2	3	4	5	9



B. The numbers behind the charts

Inbox factors (Charts 8 to 11)

Factor	Value
Verified sending (SPF, DKIM, DMARC)	108 contacts per reply
Not verified	120 contacts per reply
Gmail	90 per reply, about 1 failed email in 34
Outlook	120 per reply, about 1 failed email in 23
Custom servers	179 per reply, about 1 failed email in 75
Warmup inbox placement (higher is better)	Outlook 99.04%, custom servers 98.83%, Gmail 98.83%
Days to first campaign (typical and average)	2024 to 2025: 9.4 and 28.6. 2025 to 2026: 8.0 and 19.8

Message and timing (Charts 12, 14, 15, 16, 17), contacts per reply

Signal	Value
Email number 1 to 7	85, 137, 179, 204, 270, 286, 303
Word count, under 50 up to 200 or more	68, 74, 97, 76, 85
Best send hour (7am UTC) and the busy block	75, and about 120
Day of week, best to worst	Wed 110, Tue 111, Sun 111, Thu 112, Mon 114, Fri 115, Sat 122
Replies within 1 hour, 24 hours, 72 hours	64%, 88.5%, 93.9%





Smartlead · The State of Cold Email 2026 · Based on more than 850 million emails sent through Smartlead from Jan to June 2026.

For information only, and reflects the views of Smartlead as of publication.