sending.ac

Warmup ≠ warming

Deliverability problems you could not explain, until now.
3rd party analysis of 4 sequencers' warmup.

DATE	June 15, 2026
PREPARED BY	sending.ac Deliverability Advisors
REFERENCE	SA-WP-2026-001

The dip you couldn't explain

ACTION ITEM

Pause campaigns to Google and review your warmup.

It starts quietly. The campaigns that used to print meetings go soft. Reply rates slip, not enough to panic, just enough to notice. You check everything: domains healthy, warmup score 100, dashboard a wall of green. So you do what every operator does. You rotate in fresh domains, switch inbox provider, turn warmup back up. It helps for a week, then the slide returns. You are treating a symptom without confirming whether a bigger problem is hiding beneath.

What your dashboard cannot see

Your warmup score measures one thing: whether your warmup emails land in the inbox instead of spam. It says nothing about who is on the other end. This matters because warmup only pays off when the mailboxes on the other end sit with the same providers your leads do — otherwise you are building reputation in a place that cannot help you. And over the last two to three years, that changed. Free warmup pools filled up almost entirely with Microsoft, so your mailboxes build a spotless reputation with Outlook and close to none with Google.

100% health

Displayed on the dashboard, for your peace of mind.

0%*

Emails actually going to Google from your mailboxes.

**As low as 0% in some cases. Less than 20% in most cases.*

If your leads are on Gmail, and most cold email lists are, that gap is the whole game: you sit at a perfect 100 while too little of your warmup ever reaches Google to teach it your mailbox behaves like a real sender. **Pause your campaigns to Google** until you have confirmed your mailboxes warm properly to Google, not only Microsoft, and compare the live Google/Microsoft split of every pool at warm.ac. The rest of this report is how warmup quietly broke, the data from our own mailboxes that shows it, and what to do.

Two things to hold onto before we start. First, this is fixable: deliverability is what we do, and the steps here work whatever you send through. Second, none of this is permanent — warmup pools are living systems, composition shifts week to week, and a sequencer actively managing its pool can pull a dip back up as fast as it appeared. Our own data shows a pool can swing back just as fast as it fell — often right after a sequencer, once shown the problem, acts to correct it. That is the loop we are building: warm.ac surfaces the issue, collaborating

sequencers act on the pool, and your deliverability comes back. The job isn't to panic when the mix moves against you; it's to see the move while it still matters and stay where your deliverability holds.

How warmup actually works

To see how a mailbox can pass every check and still fail, look under the hood. Warmup is simple: your mailboxes trade friendly emails with a shared pool of other mailboxes, those get opened and replied to, and inbox providers learn your domain behaves like a real person, not a spammer. That earned trust is your sender reputation, and it decides whether your campaigns land.

Only two names matter

Roughly 90% of the world's inboxes belong to Google and Microsoft, so building reputation with them is the whole job. Four kinds of infrastructure feed the pools: Google runs on Google IPs, Outlook on Microsoft IPs, Azure also on Microsoft IPs (the part almost nobody realises), and SMTP senders, which can be thought of like leeches in the torrenting world: they take from the pool but add next to nothing to the sender reputation you actually care about. This is why much of the industry frowns on SMTP providers, and why many SMTPs face an uphill battle in most warmup pools — they get naturally treated less favorably than the reputation-building equivalents. Real warmup means reaching genuine Google and Microsoft inboxes.

A pool is only as good as the company it keeps

Here is the hinge the whole story turns on: your warmup is only as good as who else is in your pool. To land at Google you need Google mailboxes to warm against; to land at Microsoft you need Microsoft. Composition equals participation. So what matters is not the number on the dashboard, it is the mix of Google and Microsoft mailboxes beside you. And that mix is exactly what shifted.

How cold email got here

This did not happen overnight, and no single company caused it. It is the result of the whole industry making the same rational decision over the last two to three years.

A few years ago, pools were roughly balanced

Warmup pools used to be a reasonable mirror of the real inbox world, split between Google and Microsoft. Warmup worked, because warming into that mix built reputation with both

providers at once.

Then the industry moved to Azure

Over the past two years the market increasingly concluded that Azure is the most cost-efficient infrastructure available, and moved to it. We predicted this trend two years ago and built sending.ac around it; the data in this report is the industry catching up. Two years ago that would have been a thesis, easy to dismiss as our bias. Reading it today, it is no longer a thesis — it is comprehensive data confirming the trend. The detail almost everyone misses is that Azure and Outlook run on the same Microsoft IP ranges. So as more of the industry moved to Azure, more Microsoft IPs entered the shared warmup pools.

WARMUP POOL COMPOSITION, BY PARTICIPANTS (2024–2026)

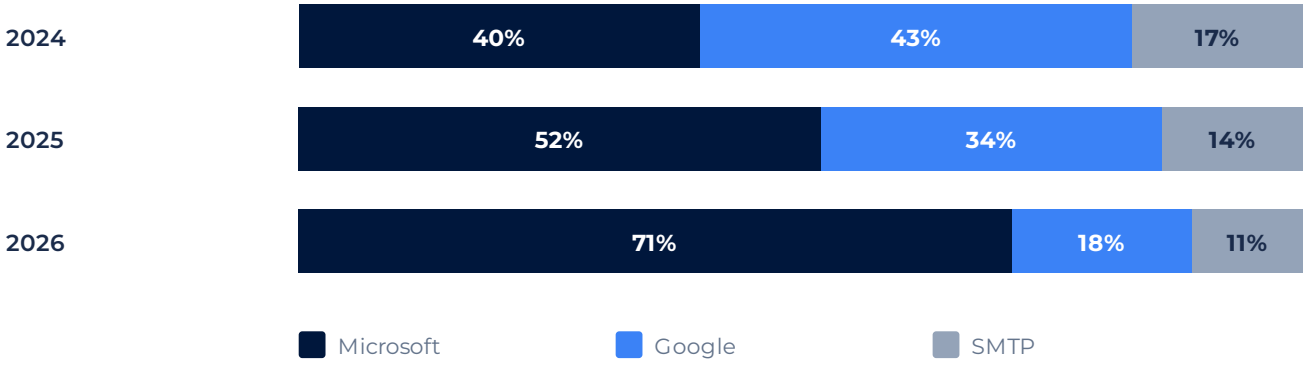

The composition tilted, and Google got crowded out

As Microsoft participation climbed, Google representation fell to low percentages. The shift is the downstream result of a collective decision, not any one operator's setup, which is exactly why it is nobody's fault and why blaming Azure misses the point. Individually switching away does not change the collective outcome of an ever-growing number of Azure users.

Now layer your lead list on top. Most cold email lists are Google-heavy, because Microsoft inboxes are harder to land in, so operators target Gmail. That is where the widening gap bites: if your warmup pool is 10% Google but your list is 90% Google, your warmup is building reputation in the wrong place.

An equal split is now mathematically impossible

Good deliverability to Google needs a pool with meaningful Google participation. Once the gap between Microsoft and Google in the pool is this wide, the balanced split your deliverability needs cannot exist, no matter how high your warmup score reads.

KEEP THIS IN MIND

Changing your infrastructure does not fix this. The problem is not what you send from, it is the destination mix of the shared pool you are in — a consequence of the entire market's choices, not yours. A Google or Outlook setup hits the same wall, because the problem is where mail is sent, not what is sending it. All of this assumes decent infrastructure: poor setups or bad sending behavior earn lesser warmup treatment in any case, to protect good senders.

The consequence

Silent deliverability dips

The result is deliverability that degrades quietly: placement at Google slips while every in-platform metric stays green, so it stays invisible until your reply rates fall. Without a third party verifying this for you, you have no way to find out.

Cold emailers fix the symptom, not the root

The usual response is to treat the symptom: switch inbox providers, replace domains, or push everything back into warmup. Sometimes those moves are justified. But none can produce good deliverability over a pool that is not warming you to the provider your leads are on. Stack enough of these and quality keeps sliding until you cross the point most operators stop tolerating.

It does not have to go this way. The same forces, without and with pool transparency:

WITHOUT TRANSPARENCY (ILLUSTRATIVE)

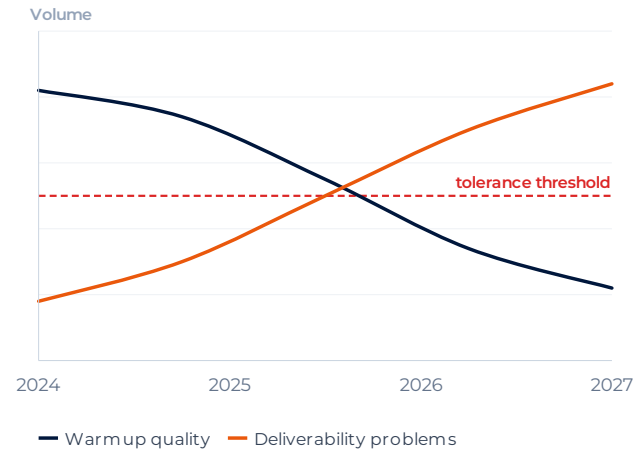

WITH TRANSPARENCY (ILLUSTRATIVE)

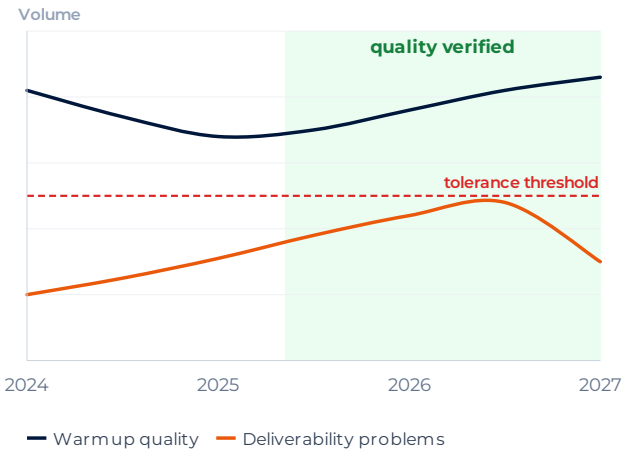

Why the industry misses it

Start with what is reasonable. Splitting users into separate pools is the right call — it shields good senders from bad ones, and the economics of keeping enough Google inboxes in a pool leave little choice in an industry that expects warmup to be free. The splitting is not the problem. The problem is that nobody tells you which pool you landed in — and that opacity is the one part of this anyone can actually fix.

Warmup is treated as a checkbox

After infrastructure, warmup quality is the biggest lever on deliverability. Yet almost nobody looks at it, because the industry treats warmup as binary, on or off, a 1 or a 0. In reality warmup quality varies enormously from one pool to the next, and composition is the first variable. A warmup score means nothing if it cannot show you the composition behind it.

Good warmup is expensive, and everyone expects it for free

Two facts that cannot both hold. A quality pool means actively recruiting and keeping Google senders in it, and that costs real money. At the same time, customers expect warmup to be free. Nobody can square that openly, so the industry has managed the gap with silence rather than by paying for quality.

So the problem persists

There is no transparency into pool composition, and the incentives are misaligned. When operators hit deliverability problems, they switched inbox providers, not sequencers, so there was never real pressure on sequencers to fix pool quality. None of this makes sequencers the villain. They are adapting to an increasingly difficult environment at Google and Outlook while customers expect warmup to cost nothing. The silence is rational. It is still a problem for you.

The proof

This began with our customers. Several reached out about deliverability slipping, we identified a cluster of affected Instantly accounts among them, and got to work. We measure every aspect of the email our infrastructure sends and receives for exactly this reason: to derive new insights from that huge stream. So we went straight to the data — SMTP-level telemetry from our own mailboxes in warmup, cross-referenced against each recipient's actual MX provider, measured outside the sequencer where its dashboards cannot see it. We present what follows as objective observations from our sample, not accusations.

Case study: a Microsoft-only pool shift we observed on Instantly (April 28)

We identified that warmup distribution for affected customers had changed, and their mailboxes were receiving virtually no Google warmup. Google share dropped from 16.9% on April 27 to near-zero — 0.3% — on April 28, with no notice, while every in-platform metric kept reading healthy.

Period	Google %	Microsoft %	What this means for you
Mar 2026 (baseline)	21–26%	53–64%	Usable, some Google warmup
Apr 19–27 (pre-event)	12–17%	70–80%	Declining, Google leaving the pool
Apr 28 – May 27	~0%	99–100%	⚠ No Google warmup whatsoever
May 28 – Jun 12 (ongoing)	Volatile, 0–99%	Volatile, 0–99%	Instantly reshuffling mailboxes between pools — composition swinging day to day

The chart below tracks pool composition day by day from April through June 2026, with the April 28 redistribution clearly visible:

INSTANTLY WARMUP POOL — GOOGLE VS MICROSOFT MIX SEEN FROM SENDING.AC CUSTOMERS, APR–JUN 2026

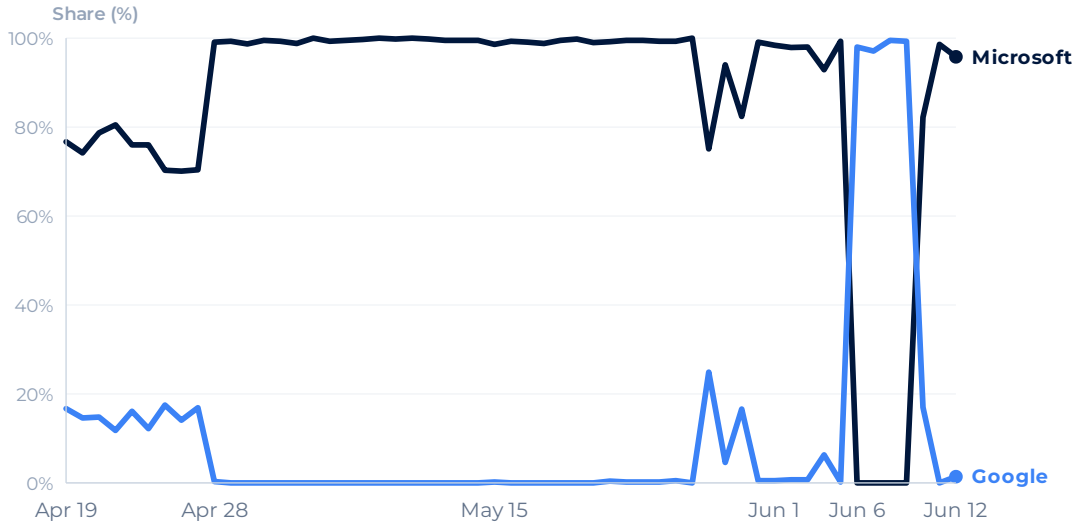

Since late May, Instantly has been reshuffling which pool each sender lands in more dynamically — separating lower-quality senders to protect pool health for everyone else. At the high volumes Instantly supports, that kind of active management is what keeps one large, low-quality sender from dragging down the mix around it. We cannot see inside

Instantly's system, so what follows is our inference from the outside, not their published criteria: judging by patterns in our own data and what is publicly known, the pool you land in does not appear fixed, and factors like reply rate, spam placement, sending volume, and personalization seem to track with the treatment a sender's pool gets. The signals we associate with warmup quality — observed, not confirmed — are described below.

Positive Signals

- High reply rates
- Personalized, differentiated outreach
- Low spam complaints from recipients
- Inbox placement on warmup emails
- OAuth-connected accounts preferred
- Sustainable volume per mailbox count

Negative Signals

- Low reply rates
- High bounce rates
- Low personalization
- High number of SMTP accounts
- Warmup emails land in spam consistently
- Loan, financial, or high-complaint offers
- Scraped or low-quality lists
- Drags down surrounding pool accounts

Complicated Signals

- Very large workspaces (100k+ mailboxes)
- High spend but low performance
- Review and solution case-by-case

Smartlead: the gap was there all along

There was no single event here. Smartlead has run separate Microsoft and SMTP pools since launch. Microsoft mailboxes carry about 48% Google, and SMTP mailboxes carry 34%. A gap of roughly 14 points, same price, invisible inside the platform. If you connect via SMTP, you have been in the weaker pool from day one.

Mailbox Type	Google %	Microsoft %	Other	What this means for you
Microsoft	~48%	~39%	~13%	✅ Better, higher Google representation
SMTP	~34%	~36%	~30%	⚠ ~14pp less Google, same price

SMARTLEAD MICROSOFT POOL — GOOGLE VS MICROSOFT MIX (MAY 25, 2026)

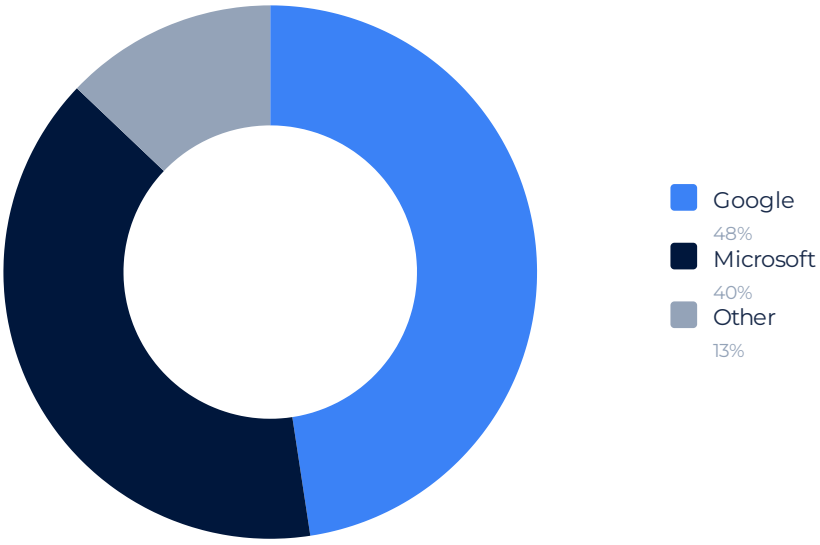

WARNING

Across every sequencer, both good and bad pools return identical in-platform metrics: warmup score, health score, inbox placement. The only way to know your real pool composition is independent telemetry measured outside the sequencer.

The industry is responding

Something changes the moment pool composition can be measured from the outside. For years warmup quality was unverifiable, so there was no pressure to compete on it. Independent measurement removes that cover: once an operator's real Google/Microsoft split is observable, the better platforms have every reason to lean into transparency rather than away from it. That shift is now underway, and it is moving in one direction.

Instantly: actively managing pool quality

The distribution shifts documented above did not go unaddressed. Instantly has moved to manage its warmup pool far more actively — reshuffling pools dynamically, separating lower-quality senders, and protecting the mix for everyone else. At the volumes Instantly supports, that kind of active management is exactly what keeps one large, low-quality sender from dragging down the pool around it. The problems became visible, and the response was to manage the pool, not just the dashboard — and active management keeps the mix in motion rather than settling it once, which is exactly why it pays to watch composition over time instead of trusting a single snapshot.

Plusvibe: opted in to independent verification

Plusvibe has opted in to having its warmup pool measured and verified by an independent third party — warm.ac. In practice that means a Plusvibe user no longer has to take pool quality on faith: the Google/Microsoft composition they warm into is confirmed by someone who does not also sell them the warmup. A quality guarantee is only worth as much as the party standing behind it, and here that party is independent.

Smartlead: opted in to independent verification

Smartlead has opted in to the same independent verification. The gap we documented earlier — where the SMTP pool carried meaningfully less Google than the Microsoft pool — is exactly the kind of thing independent measurement surfaces and keeps honest. By opting in, Smartlead gives its users a composition they can check rather than assume, validated continuously at warm.ac rather than asserted once.

EmailBison: the invitation is open

We measure every sequencer's pool the same way from the outside, whether or not they collaborate with us — EmailBison included. Opting in simply adds an independently verified quality guarantee on top of the numbers we already publish, and the door is open to every platform that wants it. The day any sequencer joins, its users gain that guarantee, and we are ready to support it the same day.

The common thread is independent verification. One by one, the platforms that compete on real deliverability are choosing to have their pools measured by a neutral party rather than vouch for themselves — because nobody should grade their own paper. warm.ac is becoming that neutral standard, and the industry is levelling up around it.

This gets solved together

There is one limit we will always be straight about. Every figure in this report is an average across a pool, and an average can never perfectly describe you. We measure a large, live sample of real mailboxes — but no outside party sees 100% of any sequencer's users, or every sub-pool they run. So the numbers we publish are a fair, honest read on the market, not a verdict on your account. We will keep making them sharper. But there is one number we can stand behind completely, and it is yours.

The one number that is truly yours

Measure your own mailboxes and the guesswork disappears. Where the aggregate tells you what is happening across the industry, your own data tells you what is happening to you —

the real Google/Microsoft split your mailboxes are actually warming into, measured directly, with no average standing in between. The aggregate proves the problem is real and widespread; only your own check proves whether it is reaching you. You need both, and the one that describes you is free to get.

And every check makes the map truer for everyone

Here is what makes this bigger than your own account. Pool composition stayed hidden for years for one reason: nobody was measuring it from the outside. That is what changes now — and it changes faster the more of us look. Every mailbox that gets measured sharpens the picture for the next operator who checks, and makes it that much harder for any pool to shift quietly without being seen. So checking your own data does two things at once: it tells you exactly where you stand, and it adds to the shared, independent map the whole industry has been missing. You get the one number that is truly yours, and you make everyone else's truer in the same move.

THE TAKEAWAY

The aggregate proves this is happening across the industry. Only your own data proves whether it is happening to you — and getting it makes the picture clearer for every operator who looks after you. That is what warm.ac is for, it is free, and it is the one number that actually describes you. Check now.

What to do

This works whatever you send through, and the first three steps cost nothing. You can run them today with data you already have.

Step 1: Find out if you're exposed (2 minutes, free)

Go to warm.ac. Find your sequencer and your mailbox type, look at the pool's Google/Microsoft split, and compare it to who you actually send to. If most of your leads are on Google but your pool is nearly all Microsoft, you are exposed: your warmup is building reputation with the wrong provider.

Step 2: Confirm it in your own data (the test that matters)

The fastest path is your own free warm.ac check, run on your mailboxes: it reports the real Google/Microsoft split your inboxes are warming into, your number rather than the pool average. Prefer to verify by hand first? Do the math yourself: take your replies, split them by each recipient's actual provider (Google versus Microsoft), and check whether Google replies

are sliding while Outlook holds steady — the signature of a Microsoft-heavy pool. It works, but it is tedious. Either way the logic is the same: the aggregate proves the problem is real across the industry, and your own data is what proves whether it is reaching you.

Step 3: Stop the bleeding (today)

1. Pause or throttle Google-recipient campaigns on affected mailboxes. Hitting real Gmail leads with zero Google warmup burns domain reputation silently, and the damage outlasts the pause.
2. Keep Microsoft-recipient campaigns running. Your Outlook deliverability is unaffected and, after a redistribution, often improved.

Step 4: Fix the pool, then re-warm

1. Measure your warmup pool composition and make sure it fits your goal. If your leads are mostly on Google, you need a pool with real Google representation — and the only way to know yours is to measure it directly, not read it off a warmup score. Where the mix doesn't fit who you send to, the answer is a better-composed pool within your existing setup, not jumping ship: measure first, then improve the pool you are already in.
2. Re-warm before you resume. Wait 2 to 4 weeks before sending to Google recipients again, and 4 weeks minimum if reply rates have already dropped.
3. Don't just switch sequencers at the same tier. The problem is the pool, not the mailboxes.

Step 5: Don't fly blind again

1. Stop treating warmup score and health score as pool-quality signals. They measure delivery, not destination.
2. Wire warm.ac into your workflow so it watches your pool continuously and flags the shift the day it happens. Pool composition can change overnight, with no notice — you want to catch the move before your reply rates do.
3. Budget for warmup. A quality pool requires active Google representation, and that costs money to maintain. If your warmup is free, you might be the product.

sending.ac measures pool composition per mailbox with independent telemetry, so we can tell you the actual Google/Microsoft split your mailboxes are warming into — not the average — and we surface the best-composed option to our customers, validated continuously at warm.ac. The steps above work whoever you buy from.

The bottom line

The trend is settled, and the data confirms it. As the industry moved to Azure, shared warmup pools filled with Microsoft and Google representation fell to a fraction of what it was. Most free pools now warm you almost entirely for Outlook — while your leads are sitting in Gmail.

It is the most expensive problem you cannot see. Your dashboard stays a wall of green while placement at Google quietly decays, and every Gmail lead you hit with too little Google warmup burns through domain reputation you cannot easily buy back. By the time your reply rates drop, the damage is already done — and it outlasts the fix.

No sequencer can solve this for you. No platform is going to be the one to tell you its own pool is failing you — nobody grades their own paper. The only way to know your real composition, and whether you have a problem at all, is an independent third party measuring from outside the platform.

And the averages are not your answer. Pools split into many sub-pools, so the aggregate is a fair read on the market, not a verdict on your mailboxes. The only data that truly describes you is your own — and getting it sharpens the picture for every operator who checks after you.

warm.ac is where you find out. Free, ongoing, and personalised to your mailboxes — it tells you whether your warmup is building reputation where your leads actually are, and keeps telling you as the pools shift underneath you.

warm.ac is who coordinates the fix. If you have a problem, we take it to your sequencer. With collaborating platforms — today Plusvibe and Smartlead — that is a direct path from diagnosis to solution; with the rest, objective measurement and a standing invitation to join.

CHECK NOW

That dip you couldn't explain? Two minutes and free, your own check tells you whether it is your pool — the one number that actually describes you. Find out where your warmup is really going, at warm.ac.

Context & method

Why this is happening, and why it isn't the sequencer's fault

Peer-to-peer warmup builds sender history by exchanging mail across a shared pool. It worked when pools were large and diverse. As cold email scaled, free pools filled with the same mailbox type and crowded out Google, just as Google and Microsoft tightened their recognition of pool behavior. A quality pool (around 40 to 50% Google) requires recruiting and retaining Google participants, which costs money. So sequencers did the rational thing and stratified: paid tiers get curated, Google-balanced pools, free tiers get whatever is left. Instantly tiered by plan, Smartlead by auth method, sorted silently. This is structural, not malicious, and it will continue through 2026.

Free warmup was quantity. Quantity is no longer sufficient. A quality pool requires active Google representation, and that costs money. If you are not paying for warmup, you are likely warming up mostly for Microsoft, barely for Google.

The question we set out to answer

Everything here answers one practical question, asked from the perspective of a new customer connecting mailboxes to a sequencer: if I connect mailboxes of type Google, Microsoft, or SMTP, and they send 100 warmup emails, how many reach Google inboxes, how many reach Microsoft, and how many reach other SMTP providers? That destination split — not the warmup score — is what decides whether your warmup builds reputation where your leads actually are.

Methodology v1.0

For the mailboxes under our management, we observe outgoing traffic and map each recipient domain to its ESP. We count warmup traffic only: a mailbox enters the warm.ac dataset for a given day only if it sends warmup that day and no campaign outreach. That isolates the pool cleanly, so the distribution reflects who you warm against, not a mix of warmup and live campaigns.

Data: SMTP-level telemetry cross-referenced with DNS-derived MX provider data, per-domain, per-day, from our own sample of mailboxes in warmup across all major sequencers, matched against an index of 6M+ tracked domains. **ESP attribution** via MX record analysis, with manual reconciliation for edge cases. **Cadence:** ongoing ingestion, index published daily.

Methodology v2.0 & future improvements

Every figure here is the observable average across a pool. Sequencers run many segregated sub-pools, often per workspace, so your individual composition may differ — some workspaces on the same tier still see healthy Google. We cannot tell which sub-pool a given account is in from this data alone.

Methodology v2.0 is about representativeness, not correctness. The measurement is valid — it reflects the warmup traffic we observe — but does not yet fully represent the average new customer joining a sequencer: weighting every mailbox equally leans the distribution toward the largest existing workspaces. We plan to weight by workspace so the numbers answer our core question directly — what an average new customer can expect, not what the highest-volume senders see.

What we are not claiming

- Not that sequencers acted maliciously. Pool stratification is a rational response to structural change.
- Not that pool composition is the only deliverability factor. Domain reputation, content, list hygiene, and sender behavior all matter.
- Not that our methodology is the only valid one. We welcome scrutiny and publish updates as it evolves.
- Not that any specific sequencer is "best". We publish data so operators can decide for themselves.

Honest incentive disclosure: our core business is sending infrastructure, so we have a commercial interest in operators upgrading. We publish this anyway because the patterns are obvious to anyone who looks, the methodology is independently verifiable, and our long-term position depends on being the most trusted voice on infrastructure, not the loudest.

TRACK IT LIVE

See live pool composition across every sequencer we track at warm.ac. Wire it into your workflow and check before every major campaign or sequencer decision. The data moves.