
# Inbox Denied
# What's Really Blocking Your B2B Email and How to Fix It


Live Q&A With ZeroBounce
August 27, 2026

## B2B EMAIL DELIVERABILITY
## Q&A GUIDE

*This document organizes the live audience questions and speaker answers from the August 27, 2026 Live Q&A, hosted by Orange Marketing and featuring ZeroBounce's Luke Glasner. Content has been edited for clarity; the content and intent are unchanged.*

# Diagnosing the Problem

## Q. How do you know if your emails are landing in the inbox or in spam?

There are a few reliable ways to check:

- Send test emails to yourself. Gather a handful of test addresses across providers and send a real campaign to them. Check whether each one lands in the primary tab, the promotions tab, or spam.
- Use an inbox placement (seed list) tool, such as ZeroBounce or similar services. These mail a "seed list" of addresses spread across many ISPs and report back where the mail landed. Placement is rarely all-or-nothing - it is usually a percentage. If a meaningful chunk of your list (roughly 20-30%) is landing in spam at a given provider, that is typically the point where it is worth investing in fixing it.
- Good tools cover business mailboxes too - Google Workspace, paid Outlook/Microsoft 365 accounts - plus solid international coverage, not just free consumer webmail.
- Watch your open-rate trend for sudden, unexplained drops. If you track opens (e.g., in HubSpot's campaign reports), the number should stay fairly consistent over time, moving only a few points based on targeting or normal day-to-day variance. A sudden drop is a signal that a portion of that campaign got routed to spam.

One added wrinkle: Apple Mail Privacy Protection (MPP), which launched in 2021, pre-fetches and "opens" every email automatically on Apple devices - but only for mail that actually reaches the inbox, since MPP does not fire from the spam folder. On lists that skew heavily toward iPhone users (one example cited: a campaign to college students), this can inflate broadcast open rates to 90%+, which is not a sign of strong engagement so much as an accidental inbox-placement test. As a rule of thumb, broadcast-style open rates consistently above 50% are worth a second look.

> *Background: engagement-metric standards were historically shaped by an industry effort called The Sane Project (circa 2020), which most ESPs adopted for calculating opens and clicks consistently. Bot activity and privacy features like Apple MPP have since added a lot of noise to that signal. ISPs like Google and Microsoft also determine engagement from in-app user behavior rather than a tracking pixel, so what they see and what your ESP reports can differ.*

## Q. If I discover a real deliverability problem, what do I do first?

Work through it in roughly this order:

- Check spam complaints first - they are the strongest negative signal your email can generate. Be careful how you read the rate: most platforms report complaints at the campaign level (e.g., 1 complaint out of 10,000 sent = 0.01%), which can mask a real problem - a provider may have only delivered 100 of those emails, one of which complained, for a true 1% complaint rate at that ISP. Look at complaints provider-by-provider, not just in aggregate.
- Sign up for Google Postmaster Tools and Yahoo Sender Hub (both free, domain-based) to see complaint volume directly, and set up spam feedback loops where available.
- Look for patterns: is a specific type of content, or a specific audience segment/list source, driving higher complaint rates? Fix that first.
- Clean the list (list hygiene). Use a list-cleansing/validation tool to remove invalid addresses and chronic "click spam" complainers before you send again. Re-validate periodically - roughly every 3 months was the recommendation given, more often for lists that have not been mailed consistently.
- Warm up your sending once the list is clean: start with your most engaged subscribers, send small, consistent volumes, and gradually increase over roughly 3-5 weeks depending on list size. This rebuilds domain/IP reputation with the mailbox providers before you return to full volume.

The reasoning behind warm-up: once a chunk of a list goes to the junk folder, those recipients stop engaging, which further lowers reputation and pushes even more mail to junk - a self-reinforcing spiral. Breaking it requires temporarily cutting volume back to very small amounts per provider, earning inbox placement again, and then stepping volume back up deliberately. Sudden, large changes in sending volume or origin (a new domain or IP suddenly sending high volume) also read as suspicious to mailbox providers - it can look like spoofing, a bot attack, or a compromised system relaying spam, and providers will throttle or block to protect both sides.

When to run a warm-up: any time you launch a new subdomain, move to a new IP without an ESP's automated warm-up process, or migrate platforms (e.g., moving a list from Mailchimp to HubSpot). It is not only a "migration" step - for many senders, warm-up is really the main lever available for actively managing domain reputation.

### Q. What should I check first from a technical standpoint?

Authentication: make sure SPF, DKIM, and DMARC are all correctly configured for your sending domain. DMARC is what records whether SPF/DKIM actually passed or failed on real mail. HubSpot will flag configuration problems directly; many other ESPs will not surface this as clearly, so it is worth checking manually. Note that a basic record lookup (e.g., a free MX Toolbox-style check) can report a record as "valid" without catching every real-world configuration issue, so do not treat a clean lookup as the final word.

## Volume, Cadence, and Context

### Q. How many emails can I send per day without getting shut down -100? 200? 1,000?

It depends on the type of sending. For 1:1 sales prospecting/sequences (e.g., through Apollo or similar tools), stay under your provider's daily threshold - Gmail, for example, caps out around 500/day. For marketing/bulk email sent through an ESP, there is no fixed daily limit.

What matters far more than the raw number is consistency. Jumping from 1,000 emails a week to 10,000 in a single send is a red flag to mailbox providers, and so is wildly inconsistent monthly volume (e.g., 10,000 one month, zero the next, 500 after that) - a common pattern seen in deliverability audits. Consistent volume and cadence protect your reputation.

- Recommended minimum cadence: at least twice a month, even though the general rule of thumb is once a month - sending twice gives you a buffer if one send gets skipped.
- Do not skip low-volume months entirely (December was called out specifically) - keep some mail flowing even during slow periods.
- Any meaningful change to your sending - a new subdomain, a big volume shift - should get its own warm-up period, since a sudden change looks unusual to mailbox providers regardless of the reason.

### Q. Do trigger words like "limited time offer," "free," or "discount" hurt deliverability?

Yes, when they read as scammy to spam filters - which is not always the same as how they read to a human recipient. It is worth testing subject lines with an AI assistant (Claude, ChatGPT) or a dedicated subject-line checker such as subjectline.com. One example shared live: a subject line using "sign up today" was flagged as sounding scammy by a spam-testing tool, even though it

read as ordinary to a person - because filters are effectively grading the line the way a machine would. "Free discount" specifically was flagged as a common triggering phrase.

### Q. Does having "marketing" (or similar terms) in your sending address increase spam risk?

Yes. The recommendation is to avoid the word "marketing" in your sending address, and to never use a "no-reply" address - regardless of what an internal IT team may suggest. Better options: something generic and human, like "hello@," or something tied to the product or brand (e.g., security@securitycompany.com for a security vendor).

Is there a "right" balance between images and text in an email, and does it vary by provider? General guidance and reasoning:

- Content-filtering rules generally favor roughly a 40/60 split of text to images, in either direction.
- Retail brands commonly send near all-image emails, but they typically slice the image into linked sections and still include alt text and some real text (e.g., in a footer).
- Reasoning: as inbox providers roll out AI-generated email summaries/previews, an all-image email leaves the AI nothing to summarize except the footer and legal boilerplate.
- Anything a recipient needs to copy and paste - a promo code, a form field value - should be real, selectable text, not baked into an image.
- All-image, no-text emails are not disqualifying on their own, but they are one of several small scoring factors that can add up with other issues into a real content-filter problem.

## Recipients & Special Sending Contexts

### Q. Do you monitor how many emails you send to a single company domain? For example, is it a red flag to send marketing and sales email to 10+ people at one company (e.g., @walmart.com)?

Not as a rule - if recipients opted in, there is no inherent cap tied to how many people at one company you are mailing. That said, large enterprises (pharma, big retailers, etc.) typically run dedicated 'messaging' or IT security teams whose job includes actively filtering and blocking marketing email from reaching employees, especially executives - independent of your sending volume.

One useful diagnostic: look up which mail platform or filtering service a domain uses. This can come from an SMTP/mail-provider field in a list-validation tool, or from a manual MX record lookup. A domain running Microsoft 365 or Google Workspace will show that directly; a domain using a filtering layer like Proofpoint or Barracuda will often show a hostname pattern (e.g., something ending in "pphosted.com" for Proofpoint) that identifies the filter in front of the mailbox.

### Q. Any special considerations for .gov or .mil recipients?

Yes - government and military recipients tend to warrant extra caution. Shared experience from working with a client with a large .gov/.mil audience, and separately from years running marketing automation for an NBA team with a similar segment: many of these environments strip links automatically, so the safest approach is plain-text email with no links, images, or buttons.

### Q. Any lessons on warming up around irregular sending patterns, like seasonal gaps?

One attendee, who ran marketing automation for an NBA team for 10 years, shared that the team had to re-warm their list every year after the off-season, ahead of new-season announcements, because the sharp drop in communications during the off-season was itself enough to trigger deliverability problems when volume ramped back up.

## List Hygiene & Engagement

### Q. How should we clean out unengaged contacts and maintain list hygiene?

List hygiene should be an ongoing habit, not a once-a-year project. If it has been more than six months since you last reviewed your list, treat that as overdue. A free list checker can help you spot-check whether a full cleaning is needed.

Historically, deliverability guidance allowed 12-18 months of inactivity before worrying about a contact. That window has shrunk: mailbox providers can recycle abandoned addresses into spam traps much faster now - reportedly in as little as ~6 months for Microsoft and roughly ~9 months for some other providers. This is part of where the common "180-day active engagement" guideline in deliverability circles comes from.

*NOTE: Spam traps and spam complaints are different things - a spam trap is a dead/ recycled address used by providers to catch poor list hygiene, while a complaint is a real recipient marking your mail as spam.*

A more precise approach than a fixed calendar rule:

- Calculate your list's average subscriber 'lifespan' (mean and median), plus roughly one standard deviation before that point - since that is typically when most drop-off happens.
- Start re-engagement before you expect to lose someone, not months after they have already gone quiet.
- Re-engagement messaging should ask directly: do you still want these emails? It can also invite feedback, or point people to a preference center (e.g., letting them choose a weekly newsletter vs. a monthly digest) so the content they get matches what they want.

Timing matters a lot here: generic, late re-engagement campaigns typically win back only 1-5% of recipients. With better-timed re-engagement (sent before full disengagement, rather than after), that can realistically reach 10-20%.

One caveat on measuring "engagement" itself: some recipients (municipal/government domains were cited as an example) block open-tracking pixels outright, so opens can understate real engagement for those segments. In those cases, look at secondary signals instead - website activity, form fills, real (non-bot) click behavior, and social engagement - and treat any activity within roughly the last 90 days as a sign the contact is still genuinely engaged.

## Q. Should we still treat open rate as a key deliverability metric, or focus on replies, clicks, and bounces instead?

Watch the trend in open rate rather than treating any single number as good or bad on its own - is it stable, did it jump, did it drop? Given how much noise bots and Apple MPP add to opens today, it is not a reliable enough signal to act on directly - for example, triggering an automated follow-up task off of an individual email open is not recommended, since an "open" increasingly may not reflect a real person.

# Testing & Strategy

## Q. What should we be A/B testing, and is our list big enough for it to matter?

Most lists are large enough to test meaningfully. For smaller lists, run the test in two waves and reverse which half gets which version the second time - that way your entire list eventually sees both versions, effectively growing your sample size.

What is worth testing: prioritize things that produce a repeatable insight, not just "which specific email copy performed better" for one send. Examples:

- One topic/angle versus another
- Short-form versus long-form copy
- A percentage-off offer versus a dollar-off offer
- Send timing - both day of week and time of day. There is a genuinely best time for your specific audience, and it will not match another company's audience. For B2B, testing around the workday is a reasonable starting point, but it is worth testing outside those hours too.

For measuring results, use conversion rate or revenue per email as the primary metric rather than opens or clicks - those are weaker proxies for the outcome you care about.

Realistically, how much does design/creative affect open and click performance versus deliverability factors like reputation and authentication? If resources are limited, what should we prioritize first? Reputation matters far more than creative. If your sender reputation has a real problem, it will hurt inbox placement more than any creative decision will help - but with a strong reputation, you get a lot more latitude on subject lines and copy. As an illustration: it is possible to intentionally use language typically considered "spammy" (profanity, all caps, words like "free") and still land reliably in the inbox, provided the sending reputation behind it is strong.

Authentication (SPF, DKIM, DMARC) is not optional - major mailbox providers will reject unauthenticated mail outright. Treat it as table stakes before investing further in creative.

### Q. What are best practices to protect your domain, and to reduce how much of this has to be handled manually?

A few structural choices cover most of it, and each one also cuts down on manual monitoring:

- Use a proper ESP so that bounce handling, spam-complaint processing, and related mechanics are managed for you automatically, and so that you get visibility into core deliverability metrics, including spam complaints. Note: many newer "outreach"-style sending tools (as distinct from traditional ESPs like Mailchimp or HubSpot) do not report spam complaints by default - if you use one of these, you will need to separately set up feedback loops and monitor tools like Google Postmaster Tools yourself.
- Use subdomains to separate different types of mail - for example, cold outreach/prospecting (higher risk) versus opt-in marketing mail (lower risk). If a subdomain used for one mail stream runs into reputation problems, it can be blocked without dragging down your other subdomains or your root domain. Without that separation, a single bad outreach campaign can jeopardize deliverability across your entire organization's email, including transactional and customer-facing mail.
- Warm up any new sending subdomain before ramping to full volume, the same as you would for a new domain or IP.

## Abuse Patterns & Emerging Trends

### Q. We worked with a company sending 2M emails/day to the same list, using a "snowshoe" strategy - buying domains in bulk and rotating to a fresh one each time one got burned. How long can an organization realistically keep this up before getting permanently burned?

This is a recognized abuse pattern, and it is actively watched for by reputation authorities - Spamhaus (spamhaus.org) specifically was named as the most consequential blacklist in this scenario. Once Spamhaus links a set of rotated domains together as belonging to the same sender, it will block the whole group. Realistically, at high daily volume, a given new domain in a rotation like this tends to get burned within about a month.

Being listed with Spamhaus specifically is far more damaging than being listed with smaller, less-influential blacklists, which often have little practical impact. The recommended habit for any serious sender: check your domains against Spamhaus regularly - weekly or monthly - rather than finding out only after deliverability collapses.

## Q. Do you believe email providers will start flagging or down-ranking content as "AI slop," similar to what platforms like LinkedIn are starting to do?

Yes, that is the expectation - providers like Google and Microsoft are likely to work signals of AI-generated content into their filtering algorithms over time, treating it as a negative signal since it does not reflect a real, engaged human sender.

What has been the most surprising email deliverability statistic over the last year? How quickly email addresses go bad - faster than expected, and especially pronounced in B2B, where job turnover has picked up noticeably. Internal data put annual address decay at roughly 28% overall, closer to 1-in-3 for B2B lists and 1-in-4 for B2C.

ZeroBounce has a new sub-status for "AI bot account" addresses - should we still market to these? The leaning shared live was to keep mailing them rather than suppress them outright, on the theory that an AI-managed inbox is still likely to summarize the message for a real human behind it - which ties back to the earlier point about avoiding all-image emails, since an AI summarizer needs text to work with.

That said, it depends on the use case: for a lower-stakes send like a newsletter or company announcement, mailing these addresses is probably fine; for something higher-stakes and narrowly targeted, like a 1:1 sales pitch, it is worth being more selective.

# Questions?

Ask us anything and stay in touch.


### Luke Glasner

**luke@zerobounce.net**
**ZeroBounce.net**
**LinkedIn Luke**

### Rebecca Gonzalez

**rebecca.gonzalez@orangemarketing.com**
**OrangeMarketing.com**
**LinkedIn Rebecca**