

The Email Authentication and Deliverability Playbook

SPF, DKIM and DMARC setup and monitoring, the Google and Yahoo bulk sender requirements, list hygiene, an honest view of warm-up, and the CAN-SPAM basics.

Contents

01 Three records, three jobs

02 Rolling out DMARC without blocking your own mail

03 What Google and Yahoo now require of bulk senders

04 List hygiene, and an honest word on warm-up

05 CAN-SPAM: what the law adds to the technical rules

06 Setup and monitoring checklist

07 What authentication will not fix

Key Takeaways & Sources

01

Three records, three jobs

Most deliverability trouble that gets blamed on "the spam filters" starts with one of three DNS records that was never set up properly, or was set up once and never revisited. SPF (RFC 7208) publishes the list of servers allowed to send mail for a domain. DKIM (RFC 6376) attaches a cryptographic signature to each message so the receiver can check that it was signed by a domain holding the matching public key, and that the signed parts were not altered in transit. DMARC (RFC 7489) sits on top of both: it asks whether the domain that passed SPF or DKIM is the same domain a person sees in the From line, and it tells receivers what to do with mail that fails, and where to send reports about it.

The alignment idea is the part people miss. A message can pass SPF for a bounce domain belonging to your email platform and still fail DMARC, because that bounce domain is not your From domain. This is why a vendor's "we handle authentication" claim is incomplete until you have checked that the platform signs with a DKIM key on your own domain, or sends from a return-path your domain controls. Most platforms offer this as a setup step involving two or three CNAME or TXT records, and it is worth doing on day one.

SPF has one hard limit that causes quiet failures: a record may trigger at most ten DNS lookups. Every include for a CRM, a ticketing tool, a marketing platform and a billing system consumes some of that budget, and an over-limit record returns a permanent error that receivers treat as a fail. The fix is to remove senders you no longer use, not to keep stacking includes. The inventory worksheet that comes with this guide exists so you can see every source in one place before you touch the record.

Record	What it proves	Common failure
SPF	The sending server's IP address is on the domain's approved list	More than ten DNS lookups, two SPF records on one domain, or a forgotten tool still listed
DKIM	The message was signed by a key the domain published, and was not altered	The platform signs with its own domain, so the signature does not align with your From address
DMARC	SPF or DKIM passed and matched the From domain; here is the policy and the reporting address	Published at p=reject on day one, or published with no one reading the reports

02

Rolling out DMARC without blocking your own mail

The safe order is to inventory first, publish a monitoring-only policy second, and tighten third. A DMARC record at `p=none` takes no action on failing mail; its value is the aggregate reports that receivers send to the address in the `rua` tag. Those reports list every IP address and source that sent mail using your domain, with pass or fail results. They are XML files and unpleasant to read raw, so most teams point the `rua` address at a DMARC reporting service or a mailbox that something parses. Whichever you pick, a named person has to look at the output on a schedule, or the record is decoration.

Expect the first reports to surprise you. Typical finds are a legacy invoicing system, a regional office using its own newsletter tool, a survey platform someone signed up for, and sometimes mail you did not send at all. Each legitimate source gets fixed so that it aligns; each illegitimate one is the reason you are doing this. Once the legitimate sources are passing consistently, move to `p=quarantine`, then to `p=reject` when you are confident. The `pct` tag lets you apply a policy to a percentage of failing mail while you build that confidence. Moving in steps is a judgement call rather than a rule from the specification, but it is the approach that avoids the worst outcome, which is a policy that silently drops a real department's mail.

Google's sender requirements ask bulk senders for a DMARC policy and allow `p=none` as the minimum. That is the compliance floor, not the goal. A domain that stays on `p=none` indefinitely is monitoring spoofing without blocking it. Whether and when to reach `p=reject` depends on how many sending sources you have and how well you control them. Mailing from a subdomain such as a dedicated marketing subdomain also lets you keep the policy on your main domain stricter than the policy on the domain where bulk mail is more likely to go wrong.

- Publish DMARC at `p=none` with a reporting address that someone actually reads
- Review reports weekly for the first month, then monthly
- Fix each legitimate source so it passes with the From domain aligned
- Move to quarantine and then reject only when legitimate sources pass consistently
- Re-check authentication whenever a new tool is allowed to send as your domain

Each legitimate source gets fixed so that it aligns; each illegitimate one is the reason you are doing this.

03

What Google and Yahoo now require of bulk senders

Google and Yahoo published matching requirements for bulk senders that took effect from February 2024. Google defines a bulk sender as one that sends close to 5,000 or more messages to personal Gmail accounts in a 24-hour period, counted across the same primary domain, and once you cross that line you stay classified as one. Yahoo's page does not state a volume cutoff. Google's published guidelines ask for SPF and DKIM authentication, a DMARC policy with the From domain aligned to SPF or DKIM, valid forward and reverse DNS for sending IPs, TLS for transmission, and a spam rate kept under limits described below. Yahoo's sender best practices set out comparable expectations. Both providers apply the full set to bulk senders and ask all senders to authenticate at least with SPF or DKIM. Read the current text rather than relying on a summary, because the providers have updated and clarified their pages since launch.

The unsubscribe rule is the one that changes day-to-day operations. For marketing and subscribed messages, bulk senders must support one-click unsubscribe as described in RFC 8058. In practice that means two headers on the message, List-Unsubscribe and List-Unsubscribe-Post, so that the mail client can show its own unsubscribe button and complete the request with a single POST, without the person logging in or confirming. Google also expects a visible unsubscribe link in the body, and recommends honouring unsubscribe requests within 48 hours. Most mainstream email platforms add these headers when you enable the setting, but custom-built sending, transactional tools that also send promotions, and sales sequencing tools are where it gets missed. Check the raw headers of a test message instead of trusting a settings page.

The spam rate threshold is the number to put on a dashboard. Google asks senders to keep the user-reported spam rate shown in Postmaster Tools below 0.10 percent and to avoid ever reaching 0.30 percent. A small sender may see little or no data in Postmaster Tools, which in our experience reflects low volume. That absence is not a clean bill of health. It means you need other signals: complaint feedback loops from the providers that offer them, unsubscribe rates, bounce rates and replies saying "why am I getting this".

For marketing and subscribed messages, bulk senders must support one-click unsubscribe as described in RFC 8058.

04

List hygiene, and an honest word on warm-up

Deliverability follows recipient behaviour more than it follows configuration. Authentication gets a message accepted as genuinely from you; it does not make anyone want it. The practical hygiene rules are unglamorous. Collect addresses with clear permission and a record of where and when it was given. Remove hard bounces immediately and permanently. Suppress anyone who unsubscribes or complains across every tool that can send to them, not only the one where they clicked. Stop sending to people who have not opened or clicked in a long time, with the window set by your sales cycle; a quarterly buyer and a weekly user are not the same audience. Never buy or rent a list: the people on it have not asked to hear from you, which is where complaint rates come from.

On warm-up, the honest position is narrow. A new domain or IP address has no sending history, and mailbox providers treat unknown senders cautiously. Increasing volume gradually, starting with your most engaged contacts and sending mail they want, gives providers a track record to judge. That is a sensible practice. What is not honest is the claim that a tool, a network of fake inboxes sending automated replies to each other, or a set number of days will produce a reputation. Providers do not publish a warm-up formula, and we have not found a primary source that does, so any specific schedule from us would be a guess. Use your own data: raise volume only while complaint, bounce and unsubscribe rates stay low, and pause when they rise.

Keep the structures that carry your reputation separate. Sales outreach, marketing newsletters and system notifications such as receipts and password resets have different complaint profiles, and a bad week in one should not drag down the others. Giving each its own subdomain is the common way to do this. It costs a little setup time and returns a clearer picture of which stream is causing trouble.

- Keep a record of consent: source, date and the wording the person saw
- Remove hard bounces on the first occurrence
- Maintain one suppression list that every sending tool respects
- Set an inactivity cut-off that matches your sales cycle, and apply it
- Separate sales, marketing and transactional streams by subdomain

A new domain or IP address has no sending history, and mailbox providers treat unknown senders cautiously.

05

CAN-SPAM: what the law adds to the technical rules

The US CAN-SPAM Act applies to any commercial email message, which includes business-to-business email; there is no exemption for sending to companies instead of consumers. The FTC's compliance guide lists the core requirements. Header information such as From, To and routing must be accurate. Subject lines must not be deceptive about the content. The message must be identifiable as an advertisement where that is not otherwise obvious. It must include a valid physical postal address for the sender. It must give a clear way to opt out, and opt-out requests must be honoured within ten business days. The opt-out cannot require a fee, ask for more personal information than an email address, or force the person through more than a single step beyond replying or visiting one web page.

Two points trip up growing teams. First, you remain responsible for mail sent on your behalf: the FTC guide states that hiring another company to handle your email does not make the legal responsibility go away. That includes agencies and sequencing vendors. Second, the law sets a minimum, and the mailbox providers set a higher practical bar. A message can meet CAN-SPAM and still be junked because the complaint rate is too high or the one-click header is missing. Treat the statute as the floor and the provider rules as the day-to-day working standard.

CAN-SPAM is a US law. Canada's anti-spam law, the EU and UK data protection and e-privacy rules, and other regimes work on different consent models, and several are stricter about prior permission. This guide does not give legal advice. If you send across borders, ask qualified counsel which regimes apply to your lists before you build the sending programme.

06

Setup and monitoring checklist

Work through the list in order. Each gate is something you can show with a test message, a DNS lookup or a screenshot, not something you decide you have done.

Step	Owner	Gate to pass before moving on
Inventory every system that sends mail using your domain	Marketing operations	Each source is on the worksheet with an owner, a purpose and a stream (sales, marketing or transactional)
Publish one SPF record and trim it to the sources you use	IT or whoever manages DNS	A single SPF record exists, within the ten-lookup limit, with no retired tools included
Enable DKIM signing on your own domain in each sending platform	Marketing operations	A test message from each platform shows a DKIM pass with the signing domain matching your From domain

Step	Owner	Gate to pass before moving on
Publish DMARC at p=none with an aggregate reporting address	IT or whoever manages DNS	Reports arrive at an address that a named person or service reads
Review reports and fix misaligned or unknown sources	Marketing operations with IT	Every legitimate source passes with alignment for two consecutive review periods
Add one-click unsubscribe headers and a visible unsubscribe link	Marketing operations	Raw headers of a test message contain List-Unsubscribe and List-Unsubscribe-Post, and the link works without a login
Connect Postmaster Tools and set a monitoring routine	Email lead	Domain verified; spam rate, bounce rate and unsubscribe rate are on a dashboard with an owner and an alert line
Tighten DMARC to quarantine, then reject, as evidence allows	IT with marketing operations	No legitimate source fails in the latest reports before each step; rollback plan written down

07

What authentication will not fix

Passing SPF, DKIM and DMARC makes your domain harder to spoof and removes a common reason for mail to be rejected. It does not make unwanted mail wanted. If recipients are marking messages as spam, the cause is the list, the offer or the frequency, and no DNS change reaches it. If you are seeing high complaint rates, look at where the addresses came from and what the message promised before touching infrastructure.

It also does not guarantee inbox placement, and nobody can honestly sell you that. Mailbox providers make their own decisions and do not publish all their signals. The best you can do is remove the avoidable reasons for rejection, keep complaints low, make unsubscribing easy and watch the numbers that are visible to you. Review the whole setup on a fixed schedule, such as quarterly, and any time a new tool is given permission to send as your domain.

Key Takeaways

- Authentication is three separate records with three separate jobs: SPF lists who may send, DKIM signs the message, and DMARC tells receivers what to do when neither lines up with the visible From domain.
- Start DMARC at p=none with reporting turned on, read the reports, fix the legitimate sources you did not know about, and only then tighten the policy.

- Google applies its strictest rules to senders of 5,000 or more messages a day to personal Gmail accounts, and Yahoo draws a similar bulk line, but authentication and a working unsubscribe are sensible for every sender.
- Complaint rate is the number to watch. Google tells senders to stay under 0.10 percent in Postmaster Tools and never reach 0.30 percent.
- There is no trick that warms a domain faster than sending wanted mail to engaged people in rising volumes. Treat any promise of instant warm-up with suspicion.

Sources

1. Google: Email sender guidelines — <https://support.google.com/a/answer/81126>
2. Google: Email sender guidelines FAQ — <https://support.google.com/a/answer/14229414>
3. Yahoo Sender Hub: Sender best practices — <https://senders.yahooinc.com/best-practices/>
4. FTC: CAN-SPAM Act, a compliance guide for business — <https://www.ftc.gov/business-guidance/resources/can-spam-act-compliance-guide-business>
5. RFC 7208: Sender Policy Framework (SPF) — <https://datatracker.ietf.org/doc/html/rfc7208>
6. RFC 6376: DomainKeys Identified Mail (DKIM) Signatures — <https://datatracker.ietf.org/doc/html/rfc6376>
7. RFC 7489: Domain-based Message Authentication, Reporting, and Conformance (DMARC) — <https://datatracker.ietf.org/doc/html/rfc7489>
8. RFC 8058: Signaling One-Click Functionality for List Email Headers — <https://datatracker.ietf.org/doc/html/rfc8058>