Messaging, Malware and Mobile Anti-Abuse Working Group
M³AAWG Recommendations for Senders Handling of Complaints

December 2017

URL to Reference this Document: www.m3aawg.org/SendersComplaintHandling

Table of Contents

I. Executive Summary

Email abuse rates can significantly affect a sender's reputation and, consequently, its ability to deliver customers' emails to the inbox. This paper explains some of the common processes senders can use to effectively manage and monitor email complaints and to help their customers, who are the list owners, develop healthy email practices that generate better results.

Understanding the impact of complaints and how to respond to them appropriately is not always clear for list owners. This lack of understanding creates a communication gap between the email technical community (e.g., postmasters, deliverability specialists, anti-spammers) and the users of email lists.

Senders receive complaints and monitor the activities of their customers to avoid receiving complaints. They usually have a technical understanding of the broader email ecosystem and a better understanding of what is acceptable, what is not acceptable, and what action is needed to mitigate problems with peers.

Senders often face challenges sharing their knowledge about complaints with their customers – the list owners. List owners often do not have the same reference framework. They are either non-technical users or marketing and sales specialists who have core competencies around marketing practices such as lead scoring, network building, business card collecting, surveying, rate of conversion, etc., and they use terms usually foreign to senders and other email specialists in their complaints.

This document describes the fundamental concepts around complaint issues in language that list owners can also recognize and appreciate. It provides examples on how to convey the right message to help list owners improve their sending practices.

II. Definitions

For the sake of clarity, below are a few concepts that should be uniquely defined. They may be the same entity but are often different business units in the same organization or are two different organizations, with one being the customer of the other.

- **List owner:** An entity with a list(s) of contacts who creates messages intended to be sent to the list.
- **Sender:** An entity that controls the messaging application, monitors delivery, mitigates issues and acts on reports.
- **Hosting company:** An entity that provides services for individuals and organizations to make their websites accessible via the World Wide Web.
- **Receiver:** A network or system that receives messages and curates incoming mail streams on behalf of one or more recipients.
- **Recipient:** The address holder to whom a message was delivered.
- **Upstream provider**: An entity that provides internet service to other internet service providers.

III. What Should be Monitored and How Often?

There are many sources of data that senders and list owners will find helpful. Some of them are external and simply require attention while others may need to be set up internally. It is important to keep the lines of communication clear from clutter as each data source is unique and issues may not be reflected across channels.

A. Looking at role accounts

Servers and computers are setup with default accounts to receive and send error reports. Left unattended, these accounts will soon be overloaded with spam, incorrectly filed requests and other useless information. Without an effective monitoring system, it will be difficult to pick up messages that identify real problems. Additionally, abuse teams should monitor all role accounts for complaints and abuse reports.

Below are the most common role addresses a good administrator will monitor for complaint issues.

1. postmaster@

Any domain that sends or receives mail must have a postmaster account. This requirement is defined in RFC5321: "Any system that includes an SMTP server supporting mail relaying or delivery MUST support the reserved mailbox 'postmaster' as a case insensitive local name."

Unfortunately, the SMTP server will also send error reports to this mailbox and, after a while, system admins tend to ignore their postmaster mailbox. It is important to fine tune the error reporting to avoid monitoring fatigue. This makes it easier to pick up a complaint or an alert sent to the postmaster.

2. hostmaster@

The Domain Names Service (DNS) specifies that every domain name must indicate an email address in the Start of Authority (SOA) record of the domain name. This is usually in the form of hostmaster@domain.name, but could be dnsadmin@domain.name, root@domain.name, hostmaster@otherdomain.name, or any similar variation.

Any address in the domain WHOIS or in DNS can also receive complaints. All listed mailboxes need to be monitored as they may receive complaints about the domain service; e.g., email, website, infrastructure, etc.

3. abuse@

The abuse mailbox is often configured and used for reporting any form of abuse coming from a domain. Commonly, the postmaster or hostmaster account is read by someone involved with the more robust infrastructure of the company but not specifically involved with anti-abuse activity. The abuse account is created for the compliance, abuse, and security specialist.

This mailbox can be indicated in the WHOIS information of the domain in the special "abuse:" field or registered on the site abuse.net, or both. Tools like spamcop.net will scrutinize these databases to find the most appropriate contact. For IP addresses, many regional internet registries now make it mandatory to have a valid and responsive abuse address to lodge problems with registered IPs.

4. Additional role accounts

There are additional role accounts that the public and upstream providers often use to report abuse.

- privacy@
- legal@
- webmaster@

5. Monitoring practices on role accounts

In addition to all the above role accounts, individuals and organizations should monitor any address that is given out or accessible to the public. This includes contact points for support and recruiting as well as any individuals with publicly listed email addresses.

Everyone in an organization who receives abuse complaints should be familiar with the abuse@ role account or another internal reporting role address. Any abuse reports received outside of that official abuse reporting alias should be forwarded to the abuse reporting alias with minimum possible delay.

6. Anti-virus and filtering on role accounts

Many role accounts are managed by the corporate mail server which has antivirus and filtering tools enabled. To ensure forwarded copies of spam get delivered, it is recommended to disable some filtering and antivirus checks for these mailboxes. To prevent confusing or discouraging users who reach out, it is recommended all incoming mail be accepted without a redirect or sending an automated message stating that no one reads the email in that inbox. This will separate good companies from spammers who sometimes create black holes or autoresponders to make people believe that the problems are addressed.

If disabling antivirus and filtering is not possible, ensure that messages are preserved through a quarantine process and that the abusive attachments are not removed. Even though abuse reports often contain or reference the abusive material directly, all such reports should be accepted and reviewed. If attachments have been removed, it is recommended that the original message with attachments be safely preserved in the case that any information within could be useful during the investigation process.

Companies can also set up a domain specifically for gathering evidence and problematic emails. Recipients reporting abuse that include evidence, such as attachments and malicious code, can be directed to forward the message to the specific domain. Using a standard domain, such as: abuse.company.example, will allow experienced reporters to bypass the normal reporting streams and provide examples directly to the appropriate department.

B. Logs

All machines and services produce logs. Important logs should be correctly identified and set up. Too much information, such as debugging data, or a lack of information may interfere with forensic investigations. In addition to the internal services which create the logs, it is important to work with vendors and service providers to ensure that monitoring and logging data from such services is reviewable and regularly assessed by both parties.

1. Raw logs

These are usually the logs produced by the mail server, and they are great for forensic work; e.g., who sent what to whom and when, was the delivery successful, etc. If these logs do not give this information, they should be fine-tuned until they do. They should retain enough history to be able to process a case that could be a few months old; for example, a recipient comes back from a vacation and complains about a two-month old spam. Specialized software can aggregate various logs and provide an interface to search through them. This software can also send alerts based on the content in the logs. A simple alert system, for example, is Logwatch.

2. High-level logs

It is likely that the application which prepares email and injects it into your mail server for final delivery also creates a log. This can often be used as a higher-level logging system that can differentiate unique customer emails and bounces, and then sorts by campaign and sending IPs, providing an overall picture of your customers' behavior and reputation.

C. Feedback

Sending emails is only one part of the equation. Getting information from receivers about the processing of these emails is also important. Senders and list owners may also receive feedback from hosting companies and upstream providers.

1. Bounces

Bounces indicate a problem with the delivery of the email and many failures are a statement on the overall reputation of the mail stream. Bounces can be a form of feedback.

The SMTP RFCs give three possible answers during a transaction:

- Success (2xy): *I have the mail*
- Temporary failure (4xy): *I cannot take this mail now, come back in a little while* - also known as a soft bounce
- Permanent failure (5xy): *I will not take this mail ever* - also known as a hard bounce

None of these replies give any indication about what to do with future mail to that address. Yet email senders must make decisions about how to handle future mail based on the three responses above.

What complicates matters is that senders have adopted the terms soft bounce and hard bounce but use them differently than above. Under the RFCs, a hard bounce (permanent failure) typically refers to a specific type of failure, where the address is nonexistent. A soft bounce typically refers to other types of failures, from rate limiting to infrastructure problems. A third type of bounce, a spam bounce, typically refers to bounces due to specific policy decisions.

a) Hard bounces

A hard bounce describes mail that could not be delivered due to a nonexistent address.

In practice, the numeric codes and associated text are not uniform and standardized across receivers. Senders and list owners should evaluate bounce codes and the text to determine why the message is rejected.

Repeatedly sending to addresses that do not exist hurts the sender's and the list owner's reputation with most ISPs. Thus, mail to addresses that do not exist should be removed from future sends. How many hard bounces are required before the address is removed from future sends is a matter of internal policy. Common numbers range from one to five hard bounces.

b) Soft bounces

A soft bounce describes mail that could not be delivered but the address is valid and future mail may be attempted. Soft bounces, in the SMTP RFC sense, are not a permanent status. Mail is retried until it either fails permanently and becomes a hard bounce or is accepted. However, senders must monitor soft bounces during a send. Significant numbers of soft bounces (temporary failures) during a send can indicate a reputation problem as many ISPs use temporary failures to slow down low-reputation mail.

In some cases, email will "soft bounce" for so long that the delivery permanently fails for that mailing. This type of soft bounce results in undelivered mail, but mail to those addresses can be attempted in the future without hurting overall reputation.

c) Spam bounces

Many ESPs (Email Service Providers) classify bounces due to reputation problems as spam bounces. While spam bounces are not a defined category, they are bounces that mention specific reputation or content problems in the bounce message. These bounces usually contain URLs linking to webpages further explaining the bounces.

Parsing logs for URLs, and then visiting the indicated websites, is a valuable way to identify specific problems. As part of their postmaster website, many of the major ISPs maintain webpages containing specific details explaining these bounces. Many ISPs use bounces to communicate problems with content, IP reputation, domain reputation, complaints or other specific issues related to the rejected message. These messages are valuable feedback from the ISPs.

In addition to SMTP-level bounces, some bounces come after the message was accepted by the receiving server. These asynchronous bounces are sent to the return path address. Most modern Mail Transfer Agents (MTAs) handle these types of bounces.

2. Opt-out or unsubscribe

Monitoring when the recipient decides to opt-out or unsubscribe is important. The links and instructions in the email must be working and must immediately notify the list owner to remove this contact from the appropriate list or lists. A proper unsubscribe mechanism is explicit and transparent from the subscriber's perspective.

The subscriber should know precisely which messages they are unsubscribing from. Ideally, they should be offered a preferences center giving them the option to unsubscribe more aggressively in cases where they are on multiple lists from a single list owner.

The actual unsubscribe mechanism procedure should consider the use case where an automated email filter visit links in the HTML on behalf of the user. While this use case is not representative of a best practice, it is frequent enough to merit consideration here. In such cases,

the automatic visit of the unsubscribe link should not cause the recipient to be unsubscribed from the mailing list. The defense against automatic link following needs to be considered against the need to avoid an unnecessary burden to the recipient that wants to unsubscribe.

To prevent accidental unsubscribes by email software and to facilitate receivers as they attempt to surface unsubscribe links for their customers, members of the email community have proposed a method to signal one-click functionality for list email headers[1]. Under these recommendations, senders can distinguish HTTPS POST actions to the unsubscribe link and treat that action as a one-click unsubscribe with no further manual intervention from the recipient.

3. Feedback Loops (FBLs)

FBLs are a mechanism that allow senders to receive emails in Abuse Report Format (ARF) from participating ISPs. When the user clicks on the "this is spam" button, an email is sent back to the sender to indicate the message was considered spam. It is important to correlate the number of reports per campaign versus the number of emails sent to participating ISPs to have an indication on how well the message was received.

Generally, the "this is spam" button is not available on a message already redirected to the spam folder as it would be redundant. In cases where the email is routed to the spam folder by the receiver, FBL reports are not sent to the sender.

D. Mailing lists

There are a few mailing lists dedicated to senders, postmasters, and the internet community that are worth watching. Conversations on the various lists cover notices of outages in networks, sending advice, and even general infrastructure discussions. It is important to monitor these lists for feedback from peers and upstream providers about specific sends or list owners.

Mailop[2] and NANOG[3] are a couple of popular and publicly accessible lists that may be worth monitoring.

E. Blocking lists

There are many blocking lists but they are not all administered or used equally. To monitor the reputation of your IPs, keep a watch on the most popular blocking lists such as Spamhaus[4], Spamcop[5], and Barracuda[6].

Often a listing in a blocking list will be triggered by emailing a spam trap and it is important to be able to link the event back to the list owner for immediate action, which is usually revoking sending rights till further investigation is completed. For more information, please refer to the Help - I'm on a Blocklist doc[7].

F. Inbox monitoring

Did the email reach the inbox or ended up in the junk folder? While this is always difficult to assess, some companies will offer helpful services. You can email a set of defined addresses as part of the campaign and they will report back if the email was received and whether it was delivered to the inbox or junk. While not 100 percent accurate, junk folder delivery could be a strong indicator that something may be wrong. This practice, called seeding, should be done

regularly to keep an eye on sender reputation and identify issues with specific ISPs. It should be kept in mind that poor inboxing may reduce abuse complaints on specific campaigns since messages already in the spam folder generally do not have the option to "mark as spam" because it would be redundant, as previously noted.

IV. We Have Data, But What Does It Mean?

Here is a list of common indicators of suspicious behavior, how to explain them to list owners, and how to suggest good and bad remedial actions that encourage list owners to follow best practices.

A. Opt-out requests

Opt-outs allow recipients to control what they want and do not want to receive. Unsubscribe links within the email should lead to a "preferences" type page. Make it easy to "unsubscribe from all" but also allow the recipient to decide which emails they do or do not want. The "preferences" page also serves as a marketing vehicle because it allows recipients to see other email products they may not be aware of and often includes the ability to opt-down into less frequent email.

Senders should also consider adding a "one click" unsubscribe link in the headers. The addition of this header will allow certain inbox providers to help their subscribers unsubscribe from trusted sources of bulk mail, which is far preferable to having those subscribers hit the "this is spam" button. Additionally, it can help reduce the incidence of recipients marking legitimate email as spam because the recipient has a hard time locating a way to unsubscribe.

A high percentage of opt-out requests may indicate a bad campaign or a stale list, but unsubscribe requests in the absence of spam complaints may simply reflect a decline in interest and engagement. If possible, ask recipients to provide an unsubscribe reason when completing the opt-out form. This data can indicate whether a particularly high opt-out percentage is healthy or unhealthy.

It is worth noting that some types of campaigns like reactivation or rebranding campaigns may generate higher than usual opt-outs. This is expected and is not a sign of underlying list collection issues.

B. Bounces: When do they matter?

It is recommended for the sender to actively monitor the bounce rates of list owners. Hard, soft and spam bounces can all tell an ESP much about the behavior of their customers.

Hard bounce rates above five percent can indicate problems with list collection or maintenance. We know receiving servers monitor the number of hard bounces and use this number as part of their overall reputation calculations. High hard bounce rates resulting from unknown users can lead directly to poor delivery.

The threshold where hard bounce rates becomes an issue can vary depending on the context and specifics of a send. For instance, when moving a list from one provider to another, the first send can often have a higher than expected bounce rate. High hard bounce rates can also indicate a purchased or rented list, and providers that ban purchased lists can address this with customers directly. Higher bounce rates can also be expected from the first send of organically collected addresses.

When a regularly mailed and properly handled list has an unexpectedly high bounce rate, it may be due to a problem at the ISP or an ISP removing old, abandoned addresses from their database. In this situation, investigation into the logs and any announcements from the ISP or industry sources can be helpful in determining whether and how this should be addressed.

Spam bounces often indicate an overall problem with reputation, either content domain or IP based. If only one client on a shared IP is having a spam bounce problem, it is likely their content or domain reputation is the problem and should be investigated further. If all clients on a shared IP are having a spam bounce problem, then the investigation is a little more challenging. Addressing the customers sending the most mail is generally a good place to start an investigation.

Soft bounces, in the SMTP sense, are a sign of MTA efficiency and tell us about the overall health of the sending platforms. However, the number of times an email fails to deliver due to soft bounces should be monitored. An address that is always timing out and failing should be considered to be hard bouncing and eventually be removed from a list.

Overall, bounces are typically a sign of some problem with the list. Monitoring bounces and the reasons given by the ISPs helps ensure the overall health of a sender.

C. Feedback loops: How do they matter?

This is a metric that must be closely monitored. Depending on the complaint rate, the feedback loop could indicate the sender is sending unsolicited mail, the recipient is receiving too many emails, or the email content is no longer relevant to the recipient. An analysis of demographic data from these recipients can provide insight into a mailing program.

In general, the number of reports to the FBL owner domain should certainly not exceed 0.1% (1 email for every 1,000 emails sent). Keep in mind that receivers may look at the raw count of abuse complaints or they may look at the count of abuse complaints as compared to the number of emails they received from the list owner or sender.

Senders have generally found the threshold of 0.1% to be a useful indicator that suggests serious problems with sending email. However, having a lower FBL abuse rate is not a guarantee that the send is without problems because ISPs calculate complaint percentages based on the number of emails delivered to their specific inboxes. Any messages that are delivered to the spam folder are not counted as "sent." Therefore, seemingly low complaint rates may not indicate a lack of problems, it may simply indicate a lack of inbox delivery. Sole reliance on the percent of complaints is not a valid indicator of good versus bad.

The important point to remember is that the sender's calculated FBL rate needs to be as close as possible to the receiver's calculated FBL rate. Some receivers host many domains, so if you do not consider the emails delivered to these domains, your rate may be higher. In contrast, if a message goes to the junk folder, it does not offer an opportunity for the recipient to click the "spam" button. Therefore, it is important to act quickly on high FBL rates as the rate may go down in future sends when increased filtering, due to a bad FBL rate and other issues, pushes more emails toward the junk folder.

FBLs are available from most of the major ISPs. Business-to-business senders, by extrapolating, can use them to indicate the condition of corporate domains.

D. Direct abuse reports

There are fewer complaints from recipients than from any other form of complaints, and consequently, they are the most significant. They often contain information about the reason for the complaint and why the normal complaint mechanisms have failed.

1. Someone wrote to me to complain, now what?

Here is a checklist of the questions that should be asked to better understand the abuse report:

- Are the unsubscribe links working properly?
- Did the recipient receive more than what he signed up for?
- How many emails did they receive and over what time frame?
- Analyze demographic makeup, recent activity, etc.
- Was there a verifiable opt-in[8] process for the complainer?

This checklist may help identify problems in the established processes, for instance unsubscribe links that do not work in some specific cases due to the complainant's environment. Reading the complaint message and understanding the complainant goes a long way to improving the sending process and supports good list management practices.

Customers should be made aware of what happens if they try to import or resend to addresses which have previously unsubscribed, followed a valid FBL, or reported direct abuse. If your system prevents those addresses from being added back to a list or sent additional mail it should be made clear to the customer, why. If the system does not have preventions in place to prevent those addresses from being sent to again, then the responsibility for educating the user falls to you. Recommendations for answering complaints can be found in Section V of this document.

2. Prioritizing complaints

All complaints are valid and should trigger a review of the list owner's practices but some types of complaints deserve additional scrutiny and consideration. For our purposes, we will look at three types of complaints in descending order of priority:

- Standard Abuse Reporting Format (ARF) complaints
- Manual complaints
- Escalated complaints with or without contractual obligations - trusted entities: blocklists, individuals, etc.

Standard ARF complaints, commonly known as FBLs, are sent by receivers when their users report that a message is spam. These have been discussed above in this document.

Manual complaints are submitted by mail recipients directly to the sender. In these cases, the mail recipient is knowledgeable enough to find the sender through the email headers or other means and is motivated to write to the sender. These complaints often come to the abuse@ alias but they can arrive through a variety of sources (see Section III above). The additional knowledge, motivation and effort that these complaints require often warrants a heightened scrutiny of the list owner's practices.

Like manual complaints, escalated complaints are submitted directly to the sender by trusted entities. Each sender must decide which entities they trust, but some examples could be members of the community known to operate blacklists and spam trap networks, members of law enforcement and even competitors.

Escalation is recommended in cases where the person reporting the issue has significant experience working in email anti-abuse since it is assumed this person can speak unequivocally about the abuse issue. For instance, complaints from an anti-spam agency should be escalated but complaints that simply copy an anti-spam agency should not. Some senders may have contractual obligations with certain entities which stipulate that complaints from those entities should be escalated and responded to in specific ways.

E. Are B2B senders the same as B2C senders?

Although B2B senders face the same rules of acceptance as B2C senders at the ISPs, B2B senders also need to be aware of the myriad types of corporate spam filters and the numerous ways that complaints can be handled. Corporate spam filters do not normally return complaints to the sender via a feedback loop device. Instead, they may block all mail coming from that domain or IP address. They may or may not send back any type of block notice in the return message. It is important for senders to review the corporate domains where it appears that the mail was accepted but then dropped prior to delivery to the recipient, i.e., there are no records of opens or clicks from any recipient at that domain.

The sender may need to contact the postmaster at the domain to obtain whitelist status at that company. Also note that complaints from corporate postmasters may be sent to the postmaster at the domain in question. The sending practices of the domain owner would then need to be reviewed.

In addition to their own filters, corporations may also use third party spam filters, for example, Barracuda, Postini, Brightmail, MessageLabs, etc. Some of these filters provide a standard way of reporting complaints.

In a business-to-consumer environment, postmasters of email recipients such as Google, Yahoo, Aol, and others would rather not block the email but instead flag the sender so that future emails are delivered to the junk folder. In a B2B environment, postmasters prefer to simply block the source of the emails.

B2B deployments are usually low in volume since they are sent to targeted, niche markets and the raw number of complaints may be low. However, the percent of complaints could be high enough to warrant a review of the list. B2B mailers build their customer lists from a larger variety of sources such as webpage registrations, newsletter sign-up pages, webinar attendees, people who downloaded white papers, industry associations, trade shows, etc. Many subscribers are added after simply making a purchase. There must be full disclosure of the company's email practices when the company asks for these email addresses. Complaints should be traced back to their original source and the source should be analyzed.

In a B2B environment, emailing everyone in a company or a department could be badly received. People will share the email among their coworkers and the list owner who emails too many people in the organization may likely be identified as a stalker rather than a marketer.

V. The Help Desk and How It Uses This Data

The M[3]AAWG document on abuse desk common practices[9]. explains how to set up a help desk, how to receive complaints and how to log them. Here we look at some specific issues.

A. Answering complaints

Failing to respond to a complainant quickly can lead to further aggravation and even further escalation of the complaint. Acknowledging that the complaint has been received helps to keep a good reputation and avoids a situation where the complainant escalates the issue to a third party. Answering within twenty-four hours, or one working day, will indicate that you act promptly on such issues. While you may not have a resolution within twenty-four hours, indicating you have started the process, and that some analysis is already in place, shows that your actions are not limited to list washing (removing the complainant email address from the list).

When the complainant is threatening legal actions, or the language carries a legal tone, it is recommended to have the sender's legal department contribute to the answer to the complainant.

It is recommended to respond to the complaint within the appropriate privacy protection laws that govern the interaction. Answering a complainant may be bound by the Terms of Services the sender has with the list owner or by the specific country or region within which the sender operates. It is important to account for local laws and the sender or list owner contract when releasing data to the complainant regarding what action was taken because of the complaint. Getting the list owner to indicate the source or acquisition method of each email address and when it was included in the list may help the sender to quickly answer the complainant by indicating how the list owner acquired the email address. In many cases the complainant may not remember which action led them to a specific list.

B. A viable Terms of Service (TOS) framework

Ultimately, both the sales staff and the abuse desk are working toward the continued success of the business. The abuse desk's job is the maintenance of an email delivery "ecosystem," and as part of this effort, it is sometimes necessary to remove a bad sender from the network. For a sales person, this could mean loss of commission or not meeting a specified sales goal, which may cause them to object to your actions. To counter these objections, or to avoid them altogether, there needs to be an understanding between the abuse desk and the salesperson that the suspension or termination of an account represents the best course of action to maintain and grow your good reputation as a sender.

Education is an important first step: It is important to clarify basic requirements and potential issues in terms that are relatable to the sales team before an account creates a potential risk to the network. This can help ease tensions about account suspension or closure and explain the potential damage that bad mailings can cause to the reputation of the sender such as:

- Cost of work to repair reputation
- Refunds to customers who complain of delivery issues resulting from the bad sender
- Cost of customers leaving due to fallout from bad senders' actions
- Added burden to support staff
- Decrease in reputation which could possibly hinder future sales

The sender should specify how it will handle cases of abuse by creating a set of rules and publishing them internally as well as in their Terms of Service. Having this rule set in a

document available to both users and staff is the best way to ensure transparency. These rules should address:

- What are the consequences of direct abuse complaints?
- What are the consequences if a customer is responsible for inclusion on a blacklist?

The sender should strive to keep the lines of communication between the abuse desk and sales open. The more informed and understanding the sales staff is about abuse desk best practices, the smoother things will go when a conversation must be had internally about an account's activity.

C. The mitigation process: Building a case towards an unhappy ending

Some customers hope that the sender has a very short memory, dealing with one case at a time. Some try to determine the specific thresholds so they can fly under the radar. It is therefore important to establish a tracking system that will record the nature of the issues, their category, their severity, and the outcome. Such a system will:

- Ensure that no issue is forgotten and all are addressed. Maintaining an issue history allows the sender to identify the areas where the list owner is weak and provide adequate training.
- Notify the list owner of each issue and provide detail on the reason behind it and its severity.
- Plot the number of cases and their severity over time and provide these reports to management when discussing a customer's case.

When educating list owners about best practices, the sender should be specific and leave as little room as possible for confusion and interpretation. Set clear requirements and expected outcomes in writing, and save these communications along with the issue history. Eventually, the list owner will either improve or the sender must stop the list owner from using their system. Providing an issue history to management and to the list owner will demonstrate a compelling reason to help end the relationship.

D. Educating the list owner on corrective actions

To mitigate complaints, list owners must be educated so they are aware of best practices.

As a first step, their contracts with their ESPs should include clauses as to what is acceptable and what is not:

- The practice of purchasing or appending email addresses must not be allowed.
- The list owners must have an up-to-date and meaningful privacy policy.
- When customers sign up for the mailing list, there should be expressed expectations of what they are signing up for and what they can expect to receive.
- There must be a verifiable and easily accessible opt-out process.
- There should be an easily accessible subscriber preference page.
- There should be a plan to re-engage customers that have not interacted.
- There should be a standard process to deactivate inactive recipients if deliverability issues arise.

As a second step, senders should have an internal process where they review the sending practices and complaint bounce percentages for each customer. There should be an ongoing dialog between the sender and the list owner on the severity of these risk factors. It is important to note that any review and reports may be generated automatically and the dialog with the user may be conducted through non-human systems. Regardless of the method, list owners are expected to ingest and respond to educational and corrective reports, such as:

- Monthly or quarterly review of deployment statistics.
- In depth look at the source of complaints - where did the name originate? Are there other recipients with similar profiles?
- Severity of blacklists.

As the final step, if complaints do not cease and other engagement metrics do not improve, the sender should engage the list owner in a discussion regarding ceasing business.

E. The cost of complaints

As early as possible, the sender should quantify the costs of every corrective action to be done with the list owner. For example, if the feedback loop rate is too high, a notification may be sent to the list owner with a follow up to ensure the problem is understood and corrected. The time spent interacting with the list owner has a value. This is often overlooked and many managers only consider the tremendous cost of an email block when determining the financial toll.

The cost of small events also needs to be factored into support for the list owner. With a good tracking system, the monthly cost of a list owner can be quantified. It will also make it easier to see if the account is "good" when you compare the value a list owner brings compared to the costs to support them.

F. You want to fire a paying customer; are you crazy?

The sender industry is one of the very rare industries where you can let a customer go despite them paying their bill on time and being a good customer in the usual sales and marketing sense. It is always difficult for someone new in this industry to understand why you may not want some potential customers as well as why you may want to get rid of some others. The reason generally is that while a customer may provide fixed revenue, they may also create costs that can sky rocket in terms of support, mitigation with ISPs, and building back a reputation within the email industry. It is therefore important early on, and in an ongoing manner, to assess the risks any customer can bring, and to understand how to mitigate these risks to protect your sending infrastructure.

All the stakeholders must be agreed on a clear, written decision path toward removing a customer before any problems arise. Such agreement will remove emotions from the decision process and make early remedial actions, like suspending sending rights, easier to enact. Communicating to the list owner where they stand at any point in this decision process will not only show that you are serious about the outcome but also indicate that, with a change of practices, a path forward to remain a customer is possible.

VI. Conclusion

Handling complaints begins with monitoring a variety of sources including publically accessible email addresses, role accounts, logs, mailing lists, and blacklists. Once complaints are received, it is important to acknowledge them within twenty-four hours, or one work day. A Terms of Service agreement provides a working framework that establishes reasonably boundaries and a common set of expectations for all involved.

With this framework in place, the person or team responding to abuse complaints can investigate complaints and respond appropriately by offering corrective steps that will prevent further abuse or having to remove the source of abuse. By staying committed to healthy complaint handling practices, senders can build a positive and sustainable reputation for themselves and for list owners.

Unfortunately, a lack of understanding exists between the email technical community (e.g., postmasters, deliverability specialists, anti-spammers) and the users of emails. This communication gap is created by an unequal technical understanding of the broad email ecosystem between senders – who receive complaints and monitor the activities of their customers to avoid receiving complaints – and the list owners, who are most often marketing and sales specialists.

This document presented the fundamental concepts around complaint issues in language that list owners can recognize and appreciate. It also provides examples on how to convey the right message to help list owners improve their sending practices.

VII. Glossary

DKIM	DomainKeys Identified Mail, a process to add an email header that allows certifying the association of an email with a specific domain and indicating the integrity of some key components in the email. The associated domain does not need to be related to any part in the email but usually is.
DNS	Domain Names Service, a service that maps domain names and hostnames to IPs and vice versa. It can also provide some additional information like SPF and DKIM records.
ESP	Email Service Provider, an entity that provides an email service.
Feedback Loop	A mechanism that some ESPs provide to selected third parties to receive information when an email is unwanted; i.e., when a user clicks on the spam button in its email application.
Hard Bounce	Mail that could not be delivered due to a non-existent address.
Hostmaster	The person or role account usually in charge of managing the DNS.
ISP	Internet Service Provider, an entity that provides internet access and usually other services or proxy services.
List owner	The entity who has a list(s) of contacts and creates the message to be sent to them.
Mailing List	A list of subscribers to a service or a group. Can also designate the software that handles the subscribers list and sending email.
MTA	Mail Transfer Agent, a system that transfers emails to another MTA or MUA.
MUA	Mail User Agent, better known as an email client.
Postmaster	The person managing the mail server.
Receiver	The entity that receives an email to either forward it to another mail server or deliver it to the final recipient.
Sender	The entity that controls the messaging application, monitors delivery, mitigates issues and acts on reports.
Soft Bounce	Mail that could not be delivered but the address is valid and future mail may be attempted.
Spam	What you do not want to receive in your mailbox.
Spam Button	A button placed on the interface of email software or web software to report that an email and similar emails are unwanted.
SPF	Sender Policy Framework, a way to indicate, via DNS, an IP range that contains the sending IPs for a domain name.

VIII. Bibliography

The reader is invited to look at the following M[3]AAWG documents and these other publications to gain a more technical and comprehensive knowledge on the issues presented in this document.

1. M[3]AAWG Senders Best Common Practices v3.0
   Defines how volume email senders can improve the deliverability of legitimate e-newsletters and permission-based e-marketing.

2. M[3]AAWG Abuse Desk Common Practices
   A summary of the most effective abuse desk best practices from M[3]AAWG service providers.

3. M[3]AAWG Code of Conduct (in Arabic) (in Chinese) (in French) (in Russian) (in Spanish)
   Outlines a voluntary set of principles for messaging system operators that discourages bulk messaging abuse of peer-to-peer messaging platforms.

4. M[3]AAWG Trust in Email Begins with Authentication
   Examines recent developments in standardized authentication mechanisms that have been tailored for use in email anti-abuse efforts. Provides background on authentication as a foundation for understanding current efforts.

5. TLS for Mail: M[3]AAWG Initial Recommendations
   Recommends three basic measures that messaging providers can implement relatively quickly to enhance the security and privacy of their users' mail.

6. M[3]AAWG Help - I'm on a Blocklist
   Specifically addresses delivery failures due to active blocks placed against a sender's IP address or domain and outlines several steps to remediate an IP or domain that has been included on a blocklist.

7. M[3]AAWG Feedback Reporting Recommendation
   Provides information on the benefits of ARF feedback reports for abuse, who should and should not send these reports, and how to receive and process ARF reports.

8. RFC 3463: Enhanced Mail System Status Codes
   Provides a set of extended status codes within mail systems for delivery status reports.

9. RFC 8058 Signaling One-Click Functionality for List Email Headers
   Describes a method for signaling a one-click function for the list-unsubscribe email header field.

IX. References

1. RFC-8058 Signaling One-Click Functionality for List Email Headers - https://datatracker.ietf.org/doc/draft-levine-mailbomb-header/

2. Mailop - http://archive.is/E1Iim

3. North American Network Operators' Group (NANOG) – Mail List Charter and Policy https://www.nanog.org/list

4. Spamhaus – https://www.spamhaus.org/

5. Spamcop - https://www.spamcop.net/

6. Barracuda - http://barracudacentral.org/rbl

7. M[3]AAWG Help – I'm on a Block List
https://www.m3aawg.org/sites/default/files/document/M3AAWG_Blocklist_Help_BP_2014-06.pdf

8. M[3]AAWG Sender Best Common Practices v3.0. Section 2.1
https://www.m3aawg.org/sites/default/files/document/M3AAWG_Senders_BCP_Ver3-2015-02.pdf

9. M[3]AAWG Abuse Desk Common Practices
https://www.m3aawg.org/sites/default/files/document/MAAWG_Abuse_Desk_Common_Practices.pdf

As with all M[3]AAWG documents that we publish, please check the M[3]AAWG website (www.m3aawg.org) for updates to this paper.