Messaging, Malware and Mobile Anti-Abuse Working Group

M³AAWG Sender Best Common Practices

Version 3.0

Updated February 2015

Note: This is a major rewrite from Version 2.0 published in October 2011 with a stronger focus on address collection and data transparency processes. It replaces both the MAAWG Sender Best Communications Practices and the Executive Summary documents included in the previous version. Both documents are now incorporated into this Version 3.0.

Abstract

This document gives an overview of the current best common practices for sending commercial electronic messaging, focusing on the technical and practical policy aspects of these operations. The audience for this M³AAWG Senders Best Common Practices ranges from delivery and compliance professionals working for an ESP (Email Servicer Provider) or a large sender to the marketing and management staff involved in the approval and deployment of these practices. This document addresses the practices related to the collection, usage and removal of recipient addresses by senders of electronic messaging. Where necessary, it links to other documents that provide more detail on a given subject.

While this document is limited to email communications, the M³AAWG Senders Committee intends to expand on other forms of electronic communications such as text messaging (i.e., SMS), instant messaging (IM), social networking-based messages and other platforms in due course. Nonetheless, many of the practices contained here also are directly applicable to these other types of messaging.

Table of Contents

1. Introduction

The M[3]AAWG Senders Committee developed these best common practices for electronic communications to support the M[3]AAWG mission of reducing messaging abuse. The goal of these practices is to promote and enhance the transparency of senders maintaining legitimate messaging so that both individual recipients and mailbox providers are more easily able to distinguish legitimate messaging from messaging abuse. This will enable mailbox providers to more effectively use their resources in the fight against messaging abuse and to better protect end-users.

M[3]AAWG categorically states that verifiably clear, conspicuous and informed opt-in subscriber consent is the best practice for messaging permission. This document outlines these and other best practices. However, at a minimum, acceptable messaging exchange best practices must be compliant with the requirements of the messaging and mailbox providers' Acceptable Use Policies (AUPs) and the applicable national and regional governments' legal and regulatory requirements (see Appendix B). Senders must adhere to these requirements to avoid industry and regulatory action and to avoid the risk of generating the need for new regulations. Senders should also consider joining relevant industry groups and adhering to related self-regulatory initiatives, such as those prescribed by other trade associations and email accreditation providers.

This document outlines industry Best Common Practices (BCPs). It is understood that some receiving networks or senders of electronic messaging may not fully implement all of these practices due to the complexity of their network infrastructures, public policy considerations, and the scalability of network platforms, but these processes do represent the consensus baseline for the industry.

2. Transparency of Intent for Address Collection

The following section outlines a set of best practices surrounding the collection and use of email addresses for bulk/commercial and transactional messaging. It focuses on the obligations that both the sending brand and Email Service Providers have to be transparent with the ultimate recipients of the messages. This section will discuss practices surrounding recipient subscription (i.e., opt-in), unsubscribe methodologies, and the handling of personal data.

2.1 Opt-in, Subscribe and Email Address Collection

The primary rule for sending email or other electronic messages, in particular bulk and marketing-related messages, is that the sender must have the explicit consent of the recipient prior to sending messages to an email address and prior to adding the recipient to any ongoing and repeated communications. Recipients are less likely to report messages as unwanted or abusive if they willingly signed up for the messages.

There are three levels to opt-in subscriptions. Each level builds on the previous level and is progressively more effective at ensuring the sender has the explicit consent of the recipient to send messages to that address:

1. **Single Opt-in**
    a. End users provide their email addresses to the sender.
    b. End users start receiving scheduled email messages related to the opt-in subscription.
    c. No notification of subscription is sent.
    d. No confirmation of consent is obtained.
2. **Single Opt-in with Notification (Better)**
    a. End users provide their email addresses to the sender.
    b. A notification/welcome message is sent to recipients informing them they will receive messages. (See Level 2 below for details)
    c. No confirmation of consent is obtained.

3. **Confirmed Opt-in (Best)**
    a. End users provide their email addresses to the sender.
    b. A confirmation message is sent to recipients informing them they must take action to complete their subscription, such as clicking a link or replying to the email. (See Level 3 below for details.)
    c. The recipients must take the prescribed action, such as clicking a link or replying to an email, to confirm they consent to receive future messages before any subscription messages are sent to them.

The following is a deeper examination of the levels listed above. These guidelines are applicable to both online and offline address collection methods. Each level builds on the previous; ie., everything from the first level should be included in the second level, even though the steps are not repeated in this paper.

Level 1 - Single Opt-in

Single opt-in should be used with extreme caution as it allows unconfirmed addresses to be added to mailing lists. This means typos or outright forgeries can be added unchecked, leading to complaints, bounces and eventual degradation of a senders reputation.

1. The recipient must knowingly give consent when the address is collected.
2. The consent must be clear and obvious. Notification of the intent to send messages to a recipient's email address should not be hidden in small type, legal jargon or on a separate page.
3. The sign up notification should clearly state the specific type of list(s) the recipient is joining. Where possible, the sender should allow the recipient to specify which lists they want to be included on or excluded from. (For instance, end users may be willing to receive update notifications on a product they have purchased, but not on related products.) The more control recipients have over the messages they receive, the less likely they are to report those messages as abusive.
4. The sender should notify recipients of the potential frequency and type of communications, and should give recipients a choice in each of these areas. Unexpected increases in the frequency of messages often leads to increased abuse complaints.
5. Email addresses should only be used for the purposes disclosed to recipients at the time of signup. If a secondary purpose is created after the address collection, such as an additional newsletter with different content, recipients should be given the opportunity to choose to opt-in to this secondary purpose. The secondary purpose must not be opt-in by default. Determining if a message constitutes a secondary purpose should be considered from the recipient's point of view. If recipients do not believe they have given permission, they are more likely to report the secondary messages as unwanted or abusive. In addition, some jurisdictions require this type of permission be collected before sending additional communications beyond the primary purpose.
6. Senders may include a visual example of the email that will be received. A visual example makes it more likely recipients will recognize the email message as a communication they requested when it arrives in their email box and they will be less likely to report it as messaging abuse.
7. At the time of address collection, a sender should consider what other types of data should be retained regarding the sign-up, such as IP address, date of collection and the website or event where the collection occurred. A large company should also keep track of which department originally collected the address and how. This information should be easily accessible should there be a requirement to prove consent to an individual, ISP (Internet Service Provider), RBL (Realtime Blackhole List) operator or regulator. Senders must also take into consideration the laws surrounding storage of Personally Identifying Information (PII) data.

Level 2 - Single Opt-in with Notification (Better)

8. A confirmation message should be sent (see 2.1, 2 above) immediately following an end user submitting their address to a list or as soon as possible (within 24 hours). These messages should adhere to the following guidelines when possible:
    a. The message should include the address submitted to the list and any other information the end user provided.
    b. The message should include notification of the frequency and type of content of the mailings.
    c. The message should use the same "from" address as future messages so recipients may add it to their address book or whitelist, if desired.
    d. If possible, servers sending confirmation messages should be segmented from regular bulk sending IPs.

Level 3 - Confirmed Opt-in (Best)

9. Confirmed opt-in (also called "double opt-in" or "closed-loop subscription") is the highest standard opt-in best practice. Recipients of the confirmation message are required to take an affirmative action to be added to the list. This prevents a typo or a maliciously submitted address from being added to ongoing mailings. The following guidelines should be followed, when possible:
    a. Notify end users at the time of address submission that they will need to check their email and take action on the confirmation message.
    b. The confirmation email should be simple and free of advertising to prevent it being identified and filtered as messaging abuse.

2.2 Implied (Implicit) Consent

Implied consent is collected when permission is inferred by a person's interaction with an organization. A common example of implied consent in email messaging is when a customer provides an email address at the point of sale but is given no notification that doing so will result in being added to ongoing mailings. The merchant then infers that the customer wants to be added to their marketing list. While this is often an accepted way to grow lists it does not work for all senders, is likely to cause abuse complaints, and is not considered a best practice. If a sender is using this method they should closely watch their complaints, opens, and clicks for this segment of the list to determine if changes should be made. One option is to add a distinct checkbox at the point of address collection to gain explicit consent, which is recommended as a best practice as noted above. If a sender choses to collect addresses via implied consent, it is a best practice to work toward gaining explicit consent prior to sending additional marketing messages via a "permission pass" or "confirmation" campaign.

It should be noted that whichever subscription method, or combination of methods, is used, it is important to keep track of how each individual address came to be on a list, including capturing originating/sign-up IPs, time and date stamp, a screen capture of the disclosure language and all relevant privacy policies in place at the time of subscription. This will help in future troubleshooting should issues arise with any particular segment of the list.

2.3 Email Append (from the M^{3}AAWG Position Statement)

The practice of email appending, also known as epending, refers to data matching exercises where a sender attempts to match a valid customer record with an email address. The owner of the email address has neither explicitly provided the address nor given consent to receive messages from the sender at this address.

Email appending is a direct violation of core M^{3}AAWG values. There are many reasons why email appending is abusive and leads to a large number of spam complaints and message rejections. In addition to complaints, email appending creates significant risks of violating privacy and anti-spam legislation. See the full text of the M^{3}AAWG position on email appending in the published document at http://www.maawg.org/sites/maawg/files/news/MAAWG_Epending_Position_2011-09.pdf

2.4 Unsubscribe

1. Senders must make the unsubscribe process as clear and easy to use as reasonably possible.
2. Senders must process all unsubscribe requests without delay to remain in compliance with all applicable laws and as a sign of respect for the recipient.
3. Senders should set expectations during the unsubscribe process detailing the specific timeframe in which the sender will process the unsubscribe request and from which list(s) or communication types the recipient is being unsubscribed. Keep in mind that the longer the timeframe for removing a recipient from a list, the more likely the user will consider the continued email abusive and complain.
4. Senders should have the capability to process email-based unsubscribe requests via the From and Reply-to addresses in the outbound email.
5. The unsubscribe link should contain everything necessary to successfully unsubscribe from the list, including:
    a. Subscriber ID
    b. Which list to unsubscribe from, if there are multiple list options
    c. Per-user-authentication-tokens, if necessary to prevent third parties from maliciously unsubscribing someone else.
6. Senders should adopt the List-Unsubscribe mechanism within the header of each message as described in RFC 2369.
    a. List-unsubscribe by "Mailto" refers to an encoded email address in an outbound message that corresponds only to that specific recipient. Any email message received at that specialized email address can be assumed to be an unsubscribe request submitted on behalf of that recipient even if it did not come from the subscribed address.
    b. List unsubscribe by "URL" involves providing a link to the unsubscribe functionality already offered by the sending platform. This link may be the same unsubscribe link typically included in most bulk mail messages. It is recommended that any personal information required for a working unsubscribe link be encoded to prevent malicious misuse of the unsubscribe function and to maximize the consumer's ease of use with the unsubscribe request.
7. Senders should determine a policy for the treatment of unsubscribes when no preference center with specific list options can be presented. This will be different for various types of businesses but it is important to decide if the unsubscribe should be valid for all mail or an individual list. It is a best practice to default to removing the recipient from all mail.
8. Senders should use easily readable text descriptions, instead of images, to accompany hyperlinks to a one-click online unsubscribe webpage.
9. Senders should also consider making offline unsubscribe mechanisms available, such as the ability to send requests to a postal mailing address or to accept unsubscribe phone calls. Any such process should be to a local (non-international) address or to a toll-free number to ensure the cost to the individual is low or nil.
10. When a subscriber is presented with a hyperlinked online subscription preference center that includes multiple subscription options, the unsubscribe option should be pre-checked by default for the user's currently subscribed lists.
11. When a provider makes new subscription offers available, returning subscribers should be presented with the new selections un-checked by default.

12. Subscribers must be able to unsubscribe without having to log in to a preference center or experience any other form of security challenge. Senders who want to offer a preference center experience may provide a login form to the subscriber so they can manage other subscriptions. However, recipients must be able to unsubscribe to a specific list outside of any secure area.

13. It is **strongly recommended** that senders include the recipient's email address in the message body to remind recipients what email address they used to subscribe to a particular list. This is particularly useful to recipients who use multiple email addresses forwarded to one central account or mailbox.

2.5 Data Security

ESPs and other senders engaged in data storage and management of subscriber email addresses or other personally identifiable information (PII) are strongly encouraged to maintain a comprehensive security program which leverages industry standard best practices to achieve security of their environment and applications. Detailed recommendations are beyond the scope of this document, but as a starting point for learning more, the Open Web Application Security Project (OWASP - https://www.owasp.org) and SANS (https://www.sans.org) provide voluminous information relating to cybersecurity and data security risks and recommendations. The Online Trust Alliance (OTA -https://otalliance.org) industry group publishes the "Data Protection and Breach Readiness Guide" which may also be of interest. Additionally, depending on the industry and application, it may be relevant to follow the recommendations in the ISO/IEC 27002:2013 code (http://www.iso.org/iso/home/store/catalogue_ics/catalogue_detail_ics.htm?csnumber=54533) found in the Access Control, Communications and Operations Management, and Information Security Incident Management code as a suitable basis for a comprehensive security program.

M[3]AAWG very **strongly recommends** that security be a primary consideration when developing processes for collecting and storing data. As criminals rob banks "because that's where the money is," it is naive to assume that cybercriminals are not similarly looking to email service providers and other online stewards of PII and confidential data as targets to acquire email address data for nefarious uses. Similarly, senders should not assume that because a service or software application stores "only" email addresses, that they are unlikely to be a target of cybercriminals. The stored data has value both to the owners (i.e., the subscribers) and to bad actors. It is essential that all members of the messaging ecosystem ensure that data is properly kept in trust and out of the hands of criminals. Additionally the U.S. Federal Trade Commission has issued a report with details on best practices for protecting consumers (http://ftc.gov/opa/2012/03/privacyframework.shtma).

3. Transparency of Data

The following section outlines various practices that senders should incorporate into their operations to be more transparent to receiving servers. Transparency is a fundamental principle for establishing trust in the email industry and applies to both the IP and domains of the sending servers and to any links included in the message body itself. When all information is out in the open, senders can be held accountable for both the good and bad behavior of their customers. When senders make themselves known through the mechanisms discussed below, receivers can more easily identify and communicate with them about problematic customers. With transparency, senders can build trust and the ability to deliver the desired email to inboxes while keeping spam out.

3.1 WHOIS Information

The WHOIS protocol is a method for system administrators to obtain contact information for IP address assignments or domain name administrators. Accurate WHOIS information for domains that assert responsibility for large volumes of email is a critical component of transparency.

Senders must maintain accurate, up-to-date WHOIS records to provide recipients and others with appropriate points of contact to help remedy abuse-related issues. This documentation must take place through WHOIS or rWHOIS. Equivalent information via other formats in addition to WHOIS/rWHOIS is also acceptable. Sub-allocations of IPs with net-blocks equivalent to or larger than /29 for IPv4 and /56 for IPv6 must be accurately

and completely documented, in keeping with ARIN (American Registry for Internet Numbers) and other RIR (Regional Internet Registry) policies; for example: https://www.arin.net/resources/request/reassignments.html.

In particular, the publication of anonymous or proxy domain registration WHOIS data circumvents this process and abridges the philosophy of transparency. There is no compelling business case for the use of intentionally vague or inaccurate WHOIS data for entities that do not intend to abuse messaging networks.

3.2 Email Authentication

Authentication supports transparency by further identifying the sender(s) of a message, while also contributing to the reduction or elimination of spoofed and forged addresses. When the sender is more easily and more accurately identified, receivers can make decisions based on mailing history and reputation for an authenticated domain. Authentication also leads to the detection of forgeries by the receiving domains that authenticate mail and therefore offers increased brand protection. Senders may chose to adopt some or all of the following authentication mechanisms based on their infrastructure and mailing configuration.

There are four main forms of authentication:

1. SPF (Sender Policy Framework) validates the sending IP against the return-path (machine name) and HELO domain.
2. Sender ID validates the sending IP against the "purported responsible address," i.e., the PRA (but this technology is no longer widely used).
3. DKIM (DomainKeys Identified Mail) uses a digital cryptographic signature which can be validated against a specific domain in the headers.
4. DMARC (Domain-based Message Authentication, Reporting & Conformance) gives email senders more control over how they would like receivers to handle unauthenticated email. It gives senders visibility into where their authenticated and unauthenticated email with their domain in the "From" line is being sent from and also provides the ability to set policy for messages that do not pass SPF and DKIM. Recipient ISPs are then able to use DMARC policies to more efficiently handle spoofed and phishing emails.

For more background on authentication, see the M^{3}AAWG paper, "Trust in Email Begins with Authentication" updated February 2015 at
https://www.m3aawg.org/sites/maawg/files/news/M3AAWG_Email_Authentication_Update-2015.pdf

The M^{3}AAWG DMARC Training Series videos presented by expert DMARC members provide an extensive course on using the technology: http://www.maawg.org/activities/training/dmarc-training-series

3.3 Technical IP Details

The following technical IP details are written exclusively with IPv4 in mind. Additional information for IPv6 will be added when there are more clearly established norms for it.

The intent of the following guidelines is to provide transparency of ownership and to clarify the responsibilities of the sending mail server to systems making reputation/filtering decisions. These guidelines will also provide an environment that is useful to system administrators investigating complaints or issues.

1. Forward DNS

   a. A name MUST be chosen that clearly identifies the responsible party's domain. This would be the provider/ESP when the IP address is shared and normally the customer/advertiser in dedicated IP situations.

      Note that in the latter situation where the domain name belongs to a customer/advertiser, it is up to the administrator of the customer/advertiser's domain to implement the forward DNS.

   b. The name must also clearly indicate that the machine is a server, rather than a generic pool space.

      Example: "server03.espname.com", not "pool-dhcp-456.espname.com"

   c. Especially in shared IP situations, there may be more than one name pointing to the same IP address, but the name chosen in (1a) above MUST be considered the "primary name."

   d. Except in dedicated IP situations, all mail servers belonging to the ESP/advertiser for the purpose of sending email to end users must use the same name or a small number of names at the "domain registry level," and use differentiating subdomains if needed. The goal is to keep the names clearly identifiable as one entity or as several sub-entities under the main one.

      Example: Do not use espname01.com, espname02.com. Instead use server1.espname.com, server2.espname.com, etc.

2. Reverse DNS

   a. The IP address of the sending server must have reverse DNS (also called PTR or IN-ADDR) configured.

   b. There must be only one reverse DNS name configured per IP address.

   c. The name must exactly match the primary name chosen as per (1a) above.

3. HELO names (names supplied by the server in HELO/EHLO commands during email transfer, see RFC5321.) These are often identical to the full hostname of the mail server itself.

   Note: HELO plays a factor in SPF authentication and this must be considered throughout this section.

   a. The HELO must be a resolvable hostname. A "[]" quoted IP address literal is not acceptable.

   b. In dedicated (unshared) IP environments, the HELO MUST exactly match the primary forward DNS name.

   c. In shared IP environments or when the IP is natted to multiple servers, the HELO name issued must match the primary name at least at the "domain registry level" of the primary forward DNS name. It may be an exact match for the whole name.

      Note: HELO names can sometimes be useful in diagnosing email problems with NATs (i.e., identifying the responsible server on a LAN), so it is recommended that each mail system or customer in a shared environment behind a "nat" have its own HELO name that is a subdomain of the domain registry level name.

      Example - whichever is more appropriate:

      "server01.espname.com", "server02.espname.com", ...
      "server01.brand1.espname.com", "server01.brand2.espname.com", ...

3.4 Shared vs. Dedicated IPs

Organizations that use an ESP or similar services to send mail in volume frequently have the option of provisioning their accounts in either a shared or a dedicated IP environment for their outbound mail. Each environment offers significant advantages and disadvantages. Both ESPs and their clients should evaluate a number of factors when determining which environment may be most appropriate and how it should be provisioned.

Defining IP Environments

A dedicated IP environment is one in which a distinct entity has exclusive use of and responsibility for the outbound mail through that environment. A dedicated environment may be comprised of one or more IPs; however, a dedicated environment can have only one entity responsible for its use.

In a shared IP environment, more than one distinct entity is assigned to an IP environment. A shared environment may be comprised of a single IP or a "pool" of IPs whose use is shared between the entities provisioned within the environment.

A Word about "Entities"

For the purposes of provisioning an IP environment, a distinct entity can be defined by various qualities. An entity can be a single company, a single brand within the company, or a single customer of an email service provider. In general, the "entity" can be defined as the party responsible for sending the message.

Determining Correct Provisioning

There are many different factors to consider when provisioning the environment for sending an entitie's email. Some of the important considerations are discussed below and are offered as general guidelines.

Dedicated Environments

There are a number of particular situations when a sender or an ESP may want, at least temporarily, to provision an entity within a dedicated environment.

The quality of the entity's mail or lists may be unknown or poorer than average. The sender or ESP may wish to isolate the mail from other entities to protect them from any ill reputational effects the mail may generate, to establish the quality of the mail, or to track any changes in the quality.

Conversely, the quality of the entity's mail or lists may be known to be of a higher-than-average quality. In this case, the sender or ESP may wish to isolate the entity from possible ill reputational effects generated by mail from other entities.

In addition, the sender or ESP may want to exert control over the volume of mail originating from their sending IP(s). Combining transactional and marketing email or mail from more than one entity can create irregular traffic volumes. Consistency of volume, whether high or low, plays an integral part in the determination of IP reputation and deliverability outcomes. This is particularly true when sending to free inbox providers, large domains and small businesses who host their mail with a large provider. It should be noted that volume consistency is not typically a useful metric for self-hosted small businesses making automated delivery decisions.

The entity may wish to be completely responsible for their mail and not have it identified with, or be attributed to, the ESP for reputation and authentication reasons. The entity may also require unique outbound MTA settings that differ from those implemented in a shared environment.

Finally, the entity may further desire certification, whitelisting or other enhanced deliverability services from a third party that requires a dedicated environment as a prerequisite for those services.

Shared Environments

There are a number of circumstances in which provisioning to a shared environment will be appropriate. As noted above, IP reputation is sensitive to changes in sending volume. Mail volumes from more than one entity can be combined in a shared environment to sustain consistent overall average sending volume from the shared environment to establish and maintain IP reputation.

Sharing an environment between more than one entity means that the reputation for the IP is also shared. By provisioning a number of entities within a shared environment, mistakes made by any single sender in the environment can dilute the overall reputational impact.

Finally, shared environments are typically less expensive to customers than a dedicated environment, which makes mailing for very small businesses economically feasible.

Considerations for Provisioning Shared Environments

Wherever possible, entities in a shared environment should have similar content and metrics, including complaint and bounce rates. Even so, extra diligence in vetting entities to be provisioned in a shared environment is recommended as there is a greater chance that the poor sending or list-building practices of one entity can "poison" the environment, impacting other entities provisioned in the environment.

It is recommended that mail from each entity within a shared environment be authenticated using DKIM, with a unique domain or subdomain as the entity asserting responsibility for the mail. This gives ISPs and recipient domains the opportunity to differentiate the individual entities responsible for the various mail streams originating from the shared environment when making automated deliverability decisions. This also permits individual entities provisioned within a shared environment to participate in DMARC. Finally, DKIM allows the entity to continue benefitting from the positive sending reputation earned in a prior environment and to carry that reputation with them in the future, should they be re-provisioned.

Best Practices for Provisioning Shared Environments

- Authenticate the domain of the ESP and the individual entity simultaneously. Admittedly, this is not an easy deployment to undertake and significant infrastructure must be in place to accomplish this goal.
- ESPs that offer both shared and dedicated environments may wish to establish criteria in advance for migration of entities from a shared to a dedicated environment. They should monitor mail from entities for the fulfillment of those criteria.

Best Practices for Provisioning Dedicated Environments

- Reverse DNS for each dedicated IP should be entity-specific rather than ESP-specific; e.g., entity.cust.esp.com. It should be reasonably simple to understand who the entity is by reading the RDNS.
- ESPs should have a well-defined process for warming IPs to be introduced to a dedicated environment and should follow it consistently.
- ESPs should help entities provisioned to dedicated environments in whitelisting the IP(s) where it is available.
- ESPs should help entities maintain a consistent average sending volume in order to accrete and maintain the best achievable sending reputation.

3.5 Vetting

ESPs that send large volumes of email on behalf of their clients are at the mercy of their worst clients' worst practices. All ESPs **must** have some type of pre-send vetting process to proactively identify malicious senders

before they mail and they **must** all have a post-send vetting process to monitor clients after they mail. A good vetting process will help the ESP determine the difference between the truly bad spammers and the customer who simply needs guidance on list hygiene. Vetting of clients is integral to maintaining a good reputation and decreasing messaging abuse. See the document MAAWG Vetting Best Common Practices (BCP) for an in depth guide to customer vetting.

3.6 Abuse/Feedback Loops (FBL) Handling

ESPs, as senders of large volumes of mail, will often be the recipients of complaints about that mail. Typically the largest source of complaint data senders receive are from automated Feedback Loops (FBLs) set up by mailbox providers, but ESPs also receive complaints directly to their abuse mailbox from recipients. ESPs must have a system to handle both FBL messages and direct complaint messages, and a process to determine what to do with these. Complaints are one of the key factors in determining if customers are in violation of a sender's Terms and Conditions or simply not behaving well. See the M[3]AAWG Feedback Reporting Recommendation for more detail. A new best practices document dedicated to abuse complaint handling will be published and available on the M[3]AAWG website in the near future.

3.7 Forwarding services

Sometimes an ESP MAY choose to set up smaller customers with addresses based on the ESP's sending domain for use in the headers of their messages. The ESP may then need to forward replies and other responses back to their customers. In this case the ESP should setup a mail forwarding service. Detailed best practices for this can be found in the document M[3]AAWG Email Forwarding Best Practices (http://www.maawg.org/system/files/news/MAAWG_Email_Forwarding_BP.pdf)

3.8 Connection/NDR (Non-Delivery Report) Handling

There is no guarantee that a sent email will be delivered to the target recipient, even if the email address is correct. There are scores of reasons why an email may not reach its target and several mechanisms by which these failures happen. Often, this will occur without the knowledge of the target recipient. A receiving email system may reject or return emails to the sending system, and the sending system must be able to receive and process these returns at the same volume and rate that they are sent. Every sender must make sure they have sufficient resources for both sending and receiving volumes of SMTP traffic.

When a receiver returns an email to the sender, it is usually returned synchronously, or "rejected" in the SMTP conversation initiated by the sender. Less frequently, the email may be returned asynchronously by an email sent to the sender's return-path address ("bounced"). This may happen hours or even days later. In practice, about 95% of all returns occur synchronously. Senders must be able to accept and identify the address which has bounced in order to process them correctly.

Regardless of the channel (synchronous or asynchronous), the email will normally be returned with a reason. The reason consists of a numeric "status code" and a "descriptive message." RFC 5321 defines basic SMTP responses (NDRs), while RFC 3463 defines the extended status codes and their meanings. Most MTA software, ISPs and mailbox providers will provide an honest and accurate "status code" that follows the RFC specifications. However, the "descriptive message" can be customized to meet the needs of the receiver and can vary greatly from one receiver to the next. Consequently they may only loosely follow the RFC. This is important because senders must be able to process, categorize and often report these returned emails based on both the returned emails' "status code" and their "descriptive message." Understanding and adjusting to a multitude of receivers returned messages is key to getting emails accepted and maintaining a good sender reputation.

There are three main classes of status-codes denoted by the first number in the code. For more detail on the SMTP response codes listed here and other codes, see the RFCs referenced above.

Examples include:

2xx – Success, message was accepted
4xx – Temporary Failure, message was not accepted
5xx – Permanent Failure, message was not accepted

The first class shown above is actually a success and does not need much elaboration. It means the message was accepted. If an email cannot be immediately accepted by the receiver or delivered to the recipient, it is returned with either a temporary or permanent failure code. A temporary failure generally means that the sender should re-queue the message and retry later. A permanent failure generally means that the sender should *not* re-queue it.

Temporary failures

As noted, temporary failures indicate that the sending server should re-queue the message and try again later. A common temporary failure may be that the receiving server is too busy with other incoming email and does not have sufficient resources to accept a given message or messages. Another cause may be that the receiver has detected unusual mailing characteristics from an IP address and has decided to stop accepting email from it for a set length of time.

These two returned messages from different receivers might have the same status code but different descriptive messages. In both cases, the sender should close the connection and open a new connection later. There exists in most commercial mailing systems, complex algorithms, queuing techniques and configurations for specifying when to retry, how frequently to retry, and when to give up. There is an expectation among receivers of very large volumes of email that the sending server will be able to make these adjustments in real time and be able to do this as a wholesale process for all emails sent to them by the same IP address. It is extremely important that senders handle temporary failures correctly and not treat them as permanent since some receivers will "tempfail" an initial message to see if the sender is following the RFCs correctly, as spammers generally will not do so.

Permanent Failures

A permanent failure indicates that the message should not be retried. The most common permanent failure is "user unknown." Additionally, many types of 5xx codes indicate a policy violation according to the descriptive text. Regardless of the descriptive text, these types of errors are always telling the sending server not to retry this message.

Handling NDRs

The numeric status codes 4xx and 5xx tell the sending mail server what to do. These numeric codes were not designed with marketers and other senders of bulk email in mind and do not clearly state whether or not the address should be removed from the mailing list or not. Senders who receive these bounces must look at the descriptive text to determine how to handle the specific address on their list. In some cases the descriptive text might help them diagnose a problem with their sending infrastructure or mailing content. It might also indicate that the sender's IP is on a blacklist and must be removed before further mail will get through to that domain. Senders of mail should attempt to diagnose and resolve delivery issues for addresses purported to be valid based on the descriptive text.

Having a process to handle NDRs is a must for any organization sending large volumes of mail. Many times a receiver will penalize a sending IP for too many messages that end up bouncing with permanent failures. If the receiver is telling the sender there is no one with that email address, it is very important this message not be retried and the email address be suppressed from future mailings. The volume of permanent failures is one indicator used by receivers to build a reputation about a sender and to make decisions about how to process future email from that sender. Large volumes of hard-bounces are too often indicative of a poorly managed registration process or an old or misapplied mailing list. It can subtract points from an IP's reputation score causing more delivery issues.

It is a recommended best practice that addresses should be removed from a list if they consistently bounce with any failure code or descriptive text over multiple consecutive campaigns. How many campaigns and over what length of time are up to the individual sender but generally it is considered a best practice to remove an address from the list if it bounces consecutively at least two times over two weeks or more. This should account for any server issues on the receiving side that could have caused an erroneous bounce.

4. Conclusion

The most important point that senders should take away from this document is that the end user and their expectations should be the highest priority. It does not matter what a sender is legally allowed to do or is granted the right to do under a privacy policy; what matters is that the recipient, their preferences and their expectations be respected. Obviously, a sender could follow all the recommendations outlined here and still have delivery issues or unhappy recipients, but strict adherence to these best practices should address the most egregious and serious issues. All senders should consult with their legal department on any choices made here to ensure they are in compliance with all necessary laws.

Appendix A - Useful Tools

Data Privacy

Some resources for more information on data privacy tools and best practices are:

- OWASP (Open Web Application Security Project), https://www.owasp.org)
- SANS Institute, https://www.sans.org
- OTA (Online Trust Alliance) "Data Protection and Breach Readiness Guide," https://otalliance.org
- ISO/IEC 27002:2013 from the Access Control, Communications and Operations Management, and Information Security Incident Management, http://www.iso.org/iso/home/store/catalogue_ics/catalogue_detail_ics.htm?csnumber=54533
- U.S. Federal Trade Commission report "Protecting Consumer Privacy in an Era of Rapid Change: Recommendations For Businesses and Policymakers," http://ftc.gov/opa/2012/03/privacyframework.shtm

Authentication

- "Trust in Email Begins with Authentication," by Dave Cocker and edited by Terry Zink, February 2015, https://www.m3aawg.org/sites/maawg/files/news/M3AAWG_Email_Authentication_Update-2015.pdf
- M[3]AAWG DMARC Training Series videos, Michael Adkins and Paul Midgen, DMARC.org: http://www.maawg.org/activities/training/dmarc-training-series

Relevant RFCs

- RFC 2369: The Use of URLs as Meta-Syntax for Core Mail List Commands and their Transport through Message Header Fields, http://tools.ietf.org/html/rfc2369
- RFC 3463: Enhanced Mail System Status Codes, http://tools.ietf.org/html/rfc3463
- RFC 5321: Simple Mail Transfer Protocol, http://tools.ietf.org/html/rfc5321
- RFC 5598: Internet Mail Architecture, https://tools.ietf.org/html/rfc5598

Other M[3]AAWG Relevant Best Practices

- Messaging Anti-Abuse Working Group Position on Email Appending, https://www.maawg.org/sites/maawg/files/news/MAAWG_Epending_Position_2011-09.pdf
- M[3]AAWG Email Forwarding Best Practices, http://www.maawg.org/system/files/news/MAAWG_Email_Forwarding_BP.pdf
- MAAWG Vetting Best Common Practices (BCP), https://www.m3aawg.org/sites/maawg/files/news/MAAWG_Vetting_BCP_2011-11.pdf

Appendix B - Legal Compliance Resources

There are a variety of laws that apply to the email industry. It is the responsibility of each company to consult their legal counsel to insure they are in compliance with all regional laws. The resources below include several repositories of laws from around the world.

- CAUCE-Cornell Spam Law Inbox Project: http://www.inboxproject.org
  A collection of materials related to anti-spam law

- Europa:
  http://europa.eu/legislation_summaries/internal_market/single_market_services/l24120_en.htm
  Summary of European Union data protection legislation in the electronic communications sector

- United States: Federal Communications Commission, Can-Spam:
  http://www.fcc.gov/encyclopedia/can-spam

- Canadian Anti-spam Law (CASL): http://fightspam.gc.ca

- CAUCE List of Official Documents Related to Canada's Anti-Spam Law
  http://www.cauce.org/2014/06/official-documents-related-to-canadas-anti-spam-law-casl.html

Appendix C - Glossary of Standard Terms

Wherever possible, these terms are derived from RFC 5598.

Access Provider – Any company or organization that provides End Users with access to the Internet. May or may not be the same entity which the End User uses as a Mailbox Provider.

Bulk/Marketing Messaging – Messages that are sent for the purposes of advertising or building the relationship between the brand/company and the recipient of the message and are not transactional. (Compare with Transactional Messaging.)

Confirmation Message – An email, sent to a new subscriber or recipient when they provide their email address to a sender. A Confirmation Message usually notifies the recipient that their email address is going to receive messages from the sender and seeks confirmation that the recipient provided the email address and wants to receive the messages. Confirmation is usually obtained by clicking a link in the message or replying to it.

Dedicated IP – A static IP address that is only used to send email on behalf of one sender/company/brand, which is responsible for the content of all the messages sent from that IP. The IP usually identifies itself in rDNS (reverse DNS) as being associated with that brand. (As compared with a Shared IP)

Dirty Mailing List – A list of email addresses where some or all of the addresses were obtained using poor acquisition and opt-in practices and/or where some or all of the addresses have not been kept up-to-date over time. For instance, this might be the result of poor handling of hard bounces, poor handling of unsubscribe requests and/or not mailing to the addresses in a year or more.

End User – A customer of a Mailbox Provider.

ESP (Email Service Provider) – A company that offers services to send email at volume on behalf of its customers; sometimes referred to as a "sender."

FBL (Feedback Loop) – A system used by a Mailbox Provider to provide qualified legitimate senders with copies of messages sent from an IP address belonging to the sender that the Mailbox Provider's end users have reported as spam. The system is provided so that senders can identify and address the problems causing the complaints.

Hard Bounce – A receiving MTA indicates that an email cannot be delivered to the recipient due to a permanent failure such as the email address no longer exists or has never existed, or the domain no longer exists or has never existed.

Mailbox Provider – A company who provides an email box to an end user. The company may or may not also provide end users with access to the Internet.

Messaging Abuse Complaint/Report – A Messaging Abuse Complaint or Report occurs when a recipient of a message complains about or reports a message as abuse. The most frequent mechanism for this is by clicking on the Spam button in a Web interface or MUA (mail user agent). However, it can also involve opening a ticket with a Mailbox Provider's support or abuse desk or sending an email complaint to the Mailbox Provider's support or abuse desk or to the sender of the message. It could also include making a phone call to the Mailbox Provider's support or abuse desk or to the sender of the message for the purpose of complaining about being sent the message.

Opt-in – The process of a recipient indicating that they wish to receive the messages from this sender.

Opt-out – The process of a recipient indicating that they do not wish to receive the messages from this sender.

Receiving MTA – The Mail Transport Agent that the Mailbox Provider is using to receive mail messages.

Sender – The sender of the email message; may refer to both the ESP who controls the Sending MTA used to send the message and also the brand or company that is responsible for the content of the message.

Sending MTA – The Mail Transport Agent that the sender is using to send the mail messages.

Shared IP – An IP that has many different senders/brands/ESP customers all mailing from it, usually simultaneously. It is usually identified with the ESP that owns the IP rather than one of the brands sending from that IP. A shared IP may also include large outbound ISP mail servers servicing retail customers.

Soft Bounce – The receiving MTA indicates that an email cannot be delivered to the recipient due to a temporary failure such as a full mailbox, a connection problem, a technical issue at the Mailbox Provider, or throttling of the connecting IP by the Mailbox Provider to slow down the rate of emails being delivered from that IP.

Transactional Messaging – Messages that are sent for the purpose of confirming a transaction between the sender and the recipient of the email, or for providing individual information about the status of the relationship between a sender and recipient. For example, this could be a bank account alert or a password change notification. (Compare with Bulk/Marketing Messaging.)