

Messaging, Malware and Mobile Anti-Abuse Working Group

M³AAWG Sender Best Common Practices

Version 4.0

Updated August 2026

Originally Published in February 2015

The reference URL for this document is m3aawg.org/senderbcp

Note: This is a major rewrite of “M³AAWG Sender Best Common Practices” version 2.0, published in October 2011, and version 3.0, published in February 2015, with a stronger focus on address collection and data transparency processes. It replaces both these earlier versions and the Executive Summary documents included in them. All of these documents are now incorporated into the present document, version 4.0, and have been updated where necessary to reflect current Sender best common practices.

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. It is intended for an audience that ranges from delivery and compliance professionals working for an ESP (Email Service 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. This document is not meant as legal advice.

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.

Abstract	1
1. Introduction	2
2. Best Practices for Infrastructure Management	3
Temporary failures	5
Permanent Failures	5
Handling NDRs	5
Domain Warming	7
Domain Warming: General Recommendations:	7
Domain Monitoring	7
Lookalike or Cousin Domain Names	8
Domain transparency, MX and A Records	8

Customer's sending domain vs ESPs subdomain	9
Consistency and alignment	9
Technical IP Details	10
IPv6	12
Bring your own IP	12
Bring your own IP: General Recommendations	12
Shared vs. Dedicated IPs	13
Defining IP Environments	13
A Word about "Entities"	13
Determining Correct Provisioning	13
2.8 Data Security	16
Protocol Security	16
Personal Data (PD)	17
3. WHOIS Information	17
4. Best practices for Address Collection and Management	17
1. Single Opt-in	18
2. Single Opt-in with Notification (Better)	18
3. Confirmed Opt-in (Best)	18
Level 1: Single Opt-in	18
Level 2: Single Opt-in with Notification (Better)	19
Level 3: Confirmed Opt-in (Best)	19
How do you protect your forms against ?	20
5. Vetting	23
6. Conclusion	23
Appendix A: Useful Tools	24
Data Privacy	24
Appendix B: Regulatory Resources	25
Appendix C: Glossary of Standard Terms	26

1. Introduction

The M³AAWG Senders Committee developed these best common practices for electronic communications to support the M³AAWG 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.

M³AAWG 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, or the scalability of network platforms, but these processes do represent the consensus for the industry best practices.

2. Best Practices for Infrastructure Management

2.1 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 should adopt some or all of the following authentication mechanisms based on their infrastructure and mailing configuration.

Detailed recommendations for authenticating outgoing mail are provided via [M³AAWG Email Authentication Recommended Best Practices](https://www.m3aawg.org/sites/default/files/m3aawg-email-authentication-recommended-best-practices-09-2020.pdf), <https://www.m3aawg.org/sites/default/files/m3aawg-email-authentication-recommended-best-practices-09-2020.pdf>.

Below we provide a short summary of the information in the document.

There are three main forms of authentication:

1. SPF (Sender Policy Framework) validates the sending IP against the Return-Path (also known as the machine name, Mail-From, SMTP, Envelope, or HELO) domain.
2. DKIM (DomainKeys Identified Mail) uses a digital cryptographic signature which can be validated against a specific domain in the headers.
3. 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.

Email Authentication: General Recommendations

- In order to comply with major mailbox provider requirements for inbound direct mail, senders and their customers must send mail that:
 - a. has an RFC5321.From domain with a published SPF record,
 - b. is DKIM signed,
 - c. has an RFC5322.From domain for which a DMARC policy exists
 - d. passes DMARC, which implies that either the RFC5321.From (SPF) domain or the DKIM signing domain aligns with the RFC5322.From domain and that SPF or DKIM

pass.

- Senders should encourage customers to align both DKIM and SPF for their mail when possible.
- Senders should validate that customer DNS records have been added properly at time of setup and then periodically going forward.
- Senders should evaluate all mail at send time to ensure it will be able to pass authentication checks and take corrective action if the message will fail as it could have a negative impact.
- Sending platforms must authenticate every message they send with DKIM and SPF with their own domain on behalf of customers who are unable to provide their own domain authentication.
- Senders should DKIM sign with both their own domain and a customer domain in order to take advantage of DKIM-based Feedback Loops (FBLs) and reporting.
- Senders should be in compliance with all RFCs.

2.2 Non-Delivery Report (NDR) 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. ESPs should make sure that detailed data regarding these NDRs can be made available to their customers. Too often these details are hidden by the ESP by only returning aggregate data such as “hard” and “soft” bounces. Customers may need access to the actual NDR message to troubleshoot delivery issues.

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, most NDRs occur synchronously but this can vary based on the composition of the list. B2B mail may have a higher rate of asynchronous bounces. Senders must be able to accept and identify the address which has bounced in order to process them correctly regardless of how the bounce came in.

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 exist 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. This is because some receivers will tempfail or greylist an initial message to see if the sender is following the RFCs correctly, as spammers generally do 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 blocklist/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 permanent failures (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 and cause 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.

2.3 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. ESPs should sign up for any and all available automated feedback loops and setup monitoring to ensure they continue to send feedback information.

See [M³AAWG Recommendations for Senders Handling of Complaints](#) for more detail.

2.4 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 set up a mail forwarding service. Detailed best practices for this can be found in the document [M³AAWG Email Forwarding Best Common Practices](#).

2.5 Domain Reputation, Warming, and Monitoring

The fully qualified domain name (FQDN) used in both email headers and body content are an integral part of a sender's reputation. With that in mind a sender should always take caution and consider sending volumes, frequencies, and domain appearances, and should continue to track these metrics with ongoing monitoring. In this section we will cover choosing a domain, warming it, and ongoing monitoring.

The reputation attached to any given domain or subdomain name can be used as a major part of the process of email filtering in tandem with other metrics like the email content itself or the reputation of the sending IP address.

The age of a domain name refers to the length of time it has been registered by the current owner. While a new domain name is not inherently suspicious, it may raise red flags if it is used to send a large volume of emails shortly after registration. Similarly, an established domain that has never previously been used to send mail suddenly starting to send large volumes can also be problematic. Both of these things can trigger spam filters and affect email deliverability, as it is a common tactic used by spammers to evade detection.

However, legitimate organizations may also register new domain names for various reasons, such as launching new products or rebranding. Monitoring and gradually increasing email sending volume from a new domain, as discussed in the section on domain warming below, can help establish a positive reputation over time.

Domain Warming

Reputation of a newly registered domain usually starts as **unknown**, which makes it similar to **bad** in that it does not yet have any sending history. A domain with history is generally considered better than if it has none; from the history, a MBP (Mailbox Provider) can deduce the quality of the behavior over time (good or bad), while no history would cause the filter to treat the incoming message with more scrutiny and caution.

Domain warming is a practice used to introduce a new domain to MBPs so they may sufficiently evaluate your messages for acceptance and placement. This will allow you to establish a reputation (either positive or negative) for sending your regular volumes of email. This process involves gradually increasing the volume of emails sent from the new domain over time. It helps prevent messages sent from the new domain from being flagged as spam solely due to sudden spikes in email volume, which can occur if a domain starts sending a large number of emails immediately after being set up. When sending from a new subdomain, to be safe, domain warming is recommended even if the organizational level domain already has an established reputation. Artificial warming practices are not recommended and may be illegal.

Domain Warming: General Recommendations

- Volume
 - Start low and slow; increase sending volume slowly. (Consider 6 weeks of warm-up as an average.)
 - Spread messages out over time. Do not send your daily batch all at once.
- Frequency
 - It's preferable to send campaigns on a regular basis until the domain is sufficiently warmed up.
- Warmup Monitoring
 - Check SMTP logs for deferrals or bounces to identify and triage in-flight delivery issues.
 - You may need to slow down sending to some or all domains if you see deferral spikes.
 - Review the quality of content and list to make sure you are sending to opt-in and engaged users.
 - Check available data points on the actions of the domain name and IPs using industry tools (see Appendix A).
 - Watch for unsubscribe and complaint rate spikes.
 - Watch for opens/clicks.
 - Large spikes may indicate security appliances giving your mail extra scrutiny.
 - Continued spikes from a security filter may require remediation.
 - Drops may indicate your mail is going to the junk folder.

Domain Monitoring

Ongoing monitoring is an important part of keeping your brand reputation high and your users safe. There are a number of industry tools that can be used to monitor the reputation of your domain and help you maintain your reputation. We have included several monitoring tools in [Appendix A - Useful Tools](#).

Lookalike or Cousin Domain Names

Organizations should avoid using cousin domains, also known as lookalike domains (domains that closely resemble their main brand domain), to mitigate the risk of phishing attacks and brand impersonation. Cousin domain names are often used by malicious actors to deceive users into believing they are interacting with a legitimate entity. Using cousin domains confuses end users and desensitizes them to be alert for phishing domains.

When a legitimate entity employs cousin domains (**example.com** and **example-promo.com**), two security challenges arise. The first pertains to end users, who are now being asked to trust variations of the primary domain. This may lead to a situation where the user is more easily fooled by other lookalike or cousin domains. The user learns to trust **example-promo.com**, and may then also trust **example-promotions.com** used by a phishing entity.

Additionally, the use of cousin domains increases the attack surface for attacks with lookalike domains. If the sender utilizes **example.com**, lookalike possibilities (**examp1e.com**) proliferate. A string that size may have a few thousand variants using typos and homographs. When a sender then deploys a cousin domain (**example-promo.com**), the sender creates an opportunity for another few thousand variants (**examp1e-pr0mo.com**) of that domain that now need to be monitored, and perhaps defensively registered, creating a cost to the brand.

Organizations can protect their brand reputation and users by proactively monitoring domain registrations, and proactively/defensively registering obvious cousin domains and not sending from them. Organizations should also make sure they are enforcing trademark rights and educating users about how to identify legitimate domains. A cousin domain should not be used as a way to get mail delivered if the sender is having reputation issues with their main organizational domain. This is what spammers do and is viewed negatively by MBPs.

Due to increased difficulty in validating ownership and legitimacy of cousin domains, MBPs do not recommend this practice. It is strongly recommended that an organization use a subdomain rather than a cousin domain when they need to differentiate parts of the business.

For info on when it may be considered acceptable to use lookalike or cousin domains, please review [M³AAWG Best Practices for Use of Look-Alike Domains for Security Purposes](#).

Domain Transparency, MX and A Records

MX Records (Mail Exchange Records): DNS records specify the mail server responsible for receiving email messages on behalf of a domain. MX records are crucial for email delivery because they direct incoming emails to the correct mail server. It is a best practice for any domain found in the 5321 or 5322.From address to have an MX record. Without properly configured MX records, email delivery may fail or be delayed. For more about MX Records, see [RFC 5321](#).

A Records: An A record maps domain names to IP addresses, indicating the location of a domain's web server. They are essential for web hosting and ensuring that visitors can access a website by entering its domain name into a web browser. Without correctly configured A records, a domain's website may become inaccessible. Some receivers will not deliver email on behalf of a sending domain if there is not a valid A record in place for that domain.

Additionally, all domains and subdomains contained in an email (headers and message body; i.e., the click-tracking domain) when entered into a browser should land on the domain owner's organizational domain or anti-abuse landing page for transparency.

Customer's Sending Domain vs. the ESP's Subdomain

It is strongly recommended that whenever possible, an ESP have their customer use their own domain or subdomain in the visible From: header line instead of the ESP's domain. This allows the actual sender to build their own reputation (be it positive or negative). Furthermore, when using the customer's own domain, it is a best practice to use a subdomain of the main domain name when sending for different purposes (e.g., **offers.mybrand.com** for marketing, **info.mybrand.com** for transactional, and so on). The subdomain name should be relevant to the type of email traffic sent to avoid confusion for both the recipient and the filtering system (which may include people-based review; words matter).

In some cases, smaller senders may not own or control the domain they are sending from and therefore are unable to do the necessary DNS configurations for authentication and other delegation setups. In these cases, it is preferable to send using the ESP's subdomain (or domain) rather than send emails with no authentication.

This configuration is not recommended for larger brands who do own their own domain. It is a fallback only for smaller entities who do not own their own domain or are unable to update the DNS. If you are not sending from a domain you control, you are not able to build a reputation on that domain for your brand.

The ESP should set up a unique, dedicated subdomain for each brand or customer who is unable to do this themselves. The subdomain should belong to a shared main domain but, in this case, should be a subdomain for each brand (e.g., **brand01.shared.esp.com**, **brand02.shared.esp.com**, and so on). While not as ideal as having customers use their own domains, it helps the MBP and anti-spam filters identify and classify the different senders on the ESP's infrastructure. ESPs may also define a shared subdomain used by all customers or broken up into pools. (e.g., **shared1.esp.com**, **shared2.esp.com**). While feasible, this is not recommended because it makes troubleshooting harder. Also, it makes it difficult to identify a source of abuse (internally or externally), as all customers sharing the subdomain are also sharing the domain's reputation. This solution is not a recommended best practice but can function as a catch-all for ESP customers who have not yet set up a domain or ESP subdomain for authentication purposes.

Consistency and Alignment

Any domain that appears in an email or SMTP transaction, including the Return-Path (technically, the [RFC5321](#).From), the visible From: (technically, the [RFC5322](#).From), the DKIM signature (as described in [RFC 6376](#)), the HELO domain, reverse DNS name of the IP, the Reply-To (as described in [RFC 5322](#)), and links within the email (especially click-tracking wrapped links) may be used by mailbox providers as a parameter by which to determine a sender's reputation, and may therefore affect whether mail is accepted or not. It is recommended to use a consistent organizational domain or subdomain throughout the email, and indeed may be required for some authentication mechanisms like DMARC.

Special care should be used when sharing domains across different senders (e.g., an ESP sharing a click-tracking domain across different customers). Because shared domains can be used as sources of reputation, if the shared domain causes mail to be blocked due to the poor sending practices of one sender, it may affect the delivery of mail from an unrelated sender. While shared domains may be used in various ways for various operational reasons, the risks of doing so should be considered. A better solution where unique domains are not possible would be to use unique subdomains per sender. Subdomains may also be used along with organizational domains. See section 2.5.2 above for more best practices around subdomains.

Some examples of fully aligned domain options follow:

Same subdomain used throughout message

Subdomain consistency	
Header field	Segmented subdomains
Return-Path	bounce@news.mybrand.com

visible From:	<great@news.mybrand.com> “News from MyBrand”
DKIM-Signature d=	news.mybrand.com
Links within body	https://news.mybrand.com/clickhere

Each address uses the organizational domain

Aligned organizational domains	
Header field	Subdomain derived from organization domain
Return-Path	bounce@bounce.mybrand.com
visible From:	<service@mybrand.com> “Your MyBrand transaction”
DKIM-Signature d=	mybrand.com
Links within body	https://click.mybrand.com/clickhere

For more information on domain usage visit [M3 AAWG Sending Domains Best Common Practices](#).

2.6 IP Address Considerations

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 but the practices below should be followed at a minimum.

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, or, 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 (*aligns with [RFC 1921](#)*)

- 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 (forward DNS) per 1. above. This is now a requirement for direct mail sent to major mailbox providers.

3. HELO names

HELO names are names supplied by the server in HELO/EHLO commands during email transfer; see [RFC 5321](#). 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 (however this is acceptable for internal MTA-to-MTA transactions).
- 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 NAT'd 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

- d. The HELO domain and matching reverse DNS domain of the sending IP should have an MX, a mail server to receive mail, and RFC 5321 and RFC 2142 required addresses (postmaster and abuse, primarily) enabled. It should have an A and a web server to redirect traffic to the main domain name used by the entity that owns that domain name.

IPv6

IPv6 is the next-gen iteration of IP addressing, using 128 bits of addressing instead of the 32 seen with IPv4. This allows for a much larger set of IPs and subnets. Utilizing IPv6 in a large-scale sending environment can provide a few benefits. Since sending via IPv6 appears to be a primary method of

tracking reputation, greater care should be taken regarding issues around domain reputation. Furthermore, other best practices should be adhered to, such as proper PTR records and valid MX records for domains being used in the messaging context.

An ESP can utilize the large amount of IPv6 space to subdivide their allocation amongst their customers. A fairly standard allocation with IPv6 is a /32 (do not confuse this with an IPv4 /32, a single IP). These /32 allocations could be further subdivided into /48 or /64, for example. Dividing an entire /32 into all /64s for customers would support four billion customers. The smallest allocation that should be given to a customer is a /64.

Where applicable, recommendations regarding reverse DNS for IPv4 should be followed as a minimum best practice.

Bring your own IP

As an alternative to utilizing IPs owned and maintained by the email service provider, some providers may allow senders to use IPs that they themselves own on the provider's platform. Reasons for doing this would include not needing to warm up new IPs on a new platform as well as retaining the ability to control and maintain their own IP inventory. As this necessitates additional technical setup on the service provider's side, it is generally not widely offered at the current time.

Bring your own IP: General Recommendations

1. Use RDAP, signed LOA (Letter of Authorization) or other means to verify ownership and control of the IP space.
2. The smallest IPv4 address range the customer may bring is /24.
3. The addresses in the IP address range must have a clean history.
4. Publish rDNS pointing to the owner's domain, not the ESP domain.
5. Sending from the current provider must be discontinued in preparation for the IP routing announcement to complete.

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 consist 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 consist 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 entity'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 benefiting 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.
- Senders should have a well-defined process for warming IPs to be introduced to a dedicated environment and should follow it consistently.
- Senders should help entities provisioned to dedicated environments in whitelisting the IP(s) where it is available.
- Senders should help entities maintain a consistent average sending volume in order to accrete and maintain the best achievable sending reputation.

2.7 Monitoring of Services

In this section we will discuss some basic monitoring that should be done on a sender’s infrastructure to detect issues related to customer breaches, account takeovers (ATO), sign-up form abuse, and IP reputation issues. This is all in addition to regular infrastructure monitoring to make sure all machines and services are up and running properly. We will cover several specific issues here but this list is not exhaustive.

ESPs are encouraged to use a combination of automated heuristics and human review to identify possible suspicious or abusive behavior. Irregular account activity can act as an indicator of potential issues. Some activity to look for:

- **Abnormal volume:** When a customer sends a higher volume than usual (more than 40%), even if it's not a result of importing an irregular list, the action can be flagged. It could be attributed to a malicious actor or simply a human error on the part of a customer, who should be educated to prevent repeating such errors in the future.

- **Abnormal list growth:** If a list grows gradually over a set period and then shows an unprecedented jump in list growth, this may also indicate a degradation of compliance and can be checked by validation of the additional addresses.
- **Abnormal bounce rates per MBP:** Global bounce rates are helpful for identifying major problems but it's important to consider specific criteria and rules for each Mailbox Provider (MBP). Unusual bounce rates per MBP can often highlight particular issues such as:
 - Authentication (SPF, DKIM, DMARC) failures
 - Routing loops
 - Missing or invalid PTR records
 - Exceeding connection limits for the receiving domain
 - IP reputation issues, or inclusion on blocklists.
- **Use of non-compliant domains:** As outlined in section 2.1, senders are flexible in choosing and adopting some or all authentication mechanisms based on their infrastructure and mailing configuration. Regardless of the chosen approach, it remains essential to establish monitoring procedures to validate whether the implemented mechanisms align with compliance standards.
- **Sign-in from unusual location:** Monitoring sign-ins from unusual locations can help identify potential unauthorized access or compromised accounts.
- **Lack of unsubscribe link:** Monitoring for the presence of an unsubscribe link is important for being in compliance with regulations like CAN-SPAM, CASL, and the ePrivacy Directive, among others. It also helps avoid deliverability issues and/or legal concerns.
- **Unusual message size:** Monitoring unusual message sizes helps detect anomalies that might indicate the presence of malware or potentially harmful attachments. This contributes to email security and protects recipients from malicious content.
- **Signup forms:** Monitoring signup forms is important to detect bot abuse.
 - Monitor subscription volume per customer to detect spikes.
 - Monitor subscription of individual addresses across multiple accounts.
- **Link redirect abuse:** Monitoring usage of click tracking/link redirects for abuse.
 - Monitor for changes in click volumes/high clicks on one recipients' links.
 - Monitor for links being flagged by anti-phishing vendors (see Appendix A for examples).
- **Content changes/abuse:**
 - Monitor for drastic changes in content/new URLs.
 - Monitor keywords for abusive content.
 - Watch for other anomalies.
 - Look for hosted images containing abusive content.

2.8 Data Security

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.

M³AAWG **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, 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.shtml>).

Protocol Security

When transmitting a message, care should be taken to use encryption during message transit. This applies to all types of message streams from every type of sender. The overhead for doing so is negligible, and it is the duty of the sending platform to do their best to protect user data (even if it seems innocuous). M³AAWG has published several documents relating to message delivery and secure transmission. The most basic steps are outlined in [TLS for Mail: M³AAWG Baseline Recommendations](#), which will provide more details on Opportunistic TLS for transmission. Opportunistic TLS still has some issues and some security weaknesses, such as vulnerability to MiTM attacks. Thus, M³AAWG has also published guidance on next-gen TLS mechanisms that exist for transmitting email, which provide a higher level of assurance that the message is being transmitted securely. That guidance discusses DANE and MTA-STS. While these are less fault-tolerant than Opportunistic TLS, there are a number of benefits. For additional details on our recommendations, visit our document [TLS for Mail: M³AAWG Initial Recommendations](#).

Remote resources are items such as images or video included in messages which are downloaded when the message is viewed, or visited as a result of clicking on a link in a message. All of these resources should be HTTPS. If your messages include non-HTTPS resources, some browsers may refuse to display mixed content (HTTP and HTTPS in the same window). Additionally, the call-to-action URLs within the message should be secured by default, using proper SSL/TLS certificates. There have been suggestions from MBPs that non-HTTPS resources could impact reputation and placement. Today, free (money-wise) certificates can be obtained from providers such as the [Let's Encrypt](#) project.

Personal Data (PD)

Personal Data (or PII, or PI) is any information that an entity may possess that could be used to trace or identify an individual. Examples include government ID number, place/date of birth, maiden name, biometric data, email address, and so on. It can also include information such as medical, financial, employment or other types of records that are linkable to an individual. There are many regional regulations (GDPR, DPA, HIPAA) that govern how PII must be handled. An organization sending messages should be aware of how PII is handled at the location where the entity for which they are sending exists, as well as the location and/or nationality of the recipient. The rules may even be different where the ESP operates. All of these need to be considered when accepting, storing, or transmitting data.

Special consideration should be given to the various interactions that recipients can trigger. Within a message, there are likely to be a number of URLs which allow the recipient to perform various actions. Some are to launch a brand site; some may be to manage preferences; others could be to unsubscribe, provide referrals, mailto: links, and so on. When actioned, URLs can be logged by network devices, or remote servers logging the HTTP request data. These URLs should obscure the email address or other PI related to the recipient. The PI should not exist in clear-text form as part of a URL, nor should it use a basic encoding which can be easily reversed, such as base64. It is far more desirable to use an encrypted version of that string, or some other UUID which has no relation to the recipient (other than a mapping at the ESP).

Furthermore, when an action URL leads to a landing page, that landing page should also obscure any PI that is shown on that page. This could be something like a**z@example.com, or “XXXX Main St.” The

landing page should not be easily discoverable. M³AAWG recommends pseudonymization to obfuscate the data and anonymize the information.

3. 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. Suballocations 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.

4. Best Practices for Address Collection and Management

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 ESPs 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. Additionally M³AAWG is working on a document describing the best practices of email list maintenance which will be published at a future date. The document will include additional steps to maintain an engaged and active email audience. This document will be updated with a link to the email list maintenance document once it is published.

4.1 Opt-in, Subscription Management 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 one; i.e., 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, creating invalid consent and leading to complaints, bounces, and eventual degradation of a sender's 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), or 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)

1. A confirmation message should be sent (see 2.a, b above) immediately following an end user's submitting their address to a list, or as soon as possible thereafter (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. Single opt-in messages that bounce should be handled as such and removed from the mailing list.
2. The message should include notification of the frequency and type of content of the mailings.
 - a. The message should use the same "from" address as future messages so recipients may add it to their address book or whitelist, if desired.
 - b. If possible, servers sending confirmation messages should be segmented from regular bulk sending IPs.

Level 3: Confirmed Opt-in (Best)

1. Confirmed opt-in (also called "double opt-in" or "closed-loop subscription") is the highest standard opt-in best practice. In addition to including all steps above (Levels 1 & 2), 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.
 - c. If a user updates their email address in your system, the new address should also be confirmed in this same way.

4.2 List Bombing

List bombing relies on a script filling out hundreds of subscription forms to one email address. This process launches confirmation messages, subscription notifications, or other transactional messages to the victim's mailbox to overwhelm their account and cause it to cease functioning for an extended period. Many times a single form is not abused in a significant way on its own, but as part of a massive subscription effort across hundreds of websites at the same time. The impact is exponential. This process makes it harder for a single company to recognize they are part of the problem.

How do you protect your forms against abuse?

Use multiple solutions to best prevent this type of abuse. Using several of these options together will build a strong form that is harder to recruit to a mail bombing attack. This is not an exhaustive list.

1. Add a CAPTCHA to the form that will not submit without it being properly completed.

2. Eliminate the ability for someone to map the form submission and then automate their way around the web page. Limit the systems/servers that can submit the form to your backend database systems.
3. Add in rate limits preventing forms from being submitted multiple times by the same IP over a period of time.
4. A sender or ESP hosting multiple forms should monitor for the same address being added to multiple different forms.
5. Providers that only service a specific geographical region might limit the form by region (showing the form only when the IP matches a set area).
6. Add blank fields to the form that only a bot might fill out, but a human would never see. For instance, make a visible email field and an invisible email field. Human eyes won't see them, but bots can read the code. Negate any with both fields filled.
7. Create a field that looks at the time stamp or generated key for the page load. If the submission time is less than a reasonable time you think it might take for someone to fill it out, or if that time is missing entirely, toss the submission. Some email address validation services might catch this type of behavior.
8. Change the names of the fields to words other than standard code. Instead of "firstname, lastname, email," use something like "First_Banana, Last_Apple, Em_Orange." While it might feel silly, the scripts running these submissions are looking for common field name variations to submit to. Unusual form fields won't likely register in the script.
9. Change the location of the page if it's been previously attacked. (e.g. www.example.com/signup becomes www.example.com/signup1.)
10. Consider adding a Form-Sub header for receiving mailbox providers.
<https://datatracker.ietf.org/doc/html/draft-levine-mailbomb-header-02>

4.3 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 and is likely to cause abuse complaints. It is not considered a best practice and should be avoided where possible. 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 unchecked box at the point of address collection to gain explicit consent, which is recommended as a best practice as noted above. If a sender chooses 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 stamps, 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, as data retention laws vary by jurisdiction.

4.4 Email Append (from the M³AAWG Position Statement)

The practice of email appending, also known as e-pending, 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 M³AAWG values as detailed in [M³AAWG Position on Email](#)

[Appending](#). There are many reasons why email appending is abusive. The practice 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.

4.5 Cold Email

M³AAWG defines cold email as an unsolicited email from an otherwise legitimate entity that tries to create a business relationship with a customer who has no prior relationship with the sender or business, and views such email as abusive. The [Messaging, Malware and Mobile Anti-Abuse Working Group \(M³AAWG\) Position on Cold Email](#) document provides the full details.

4.6 Unsubscribe: General Recommendations

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 one-click functionality link to the unsubscribe 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 in order to prevent malicious misuse of the unsubscribe function and to maximize the consumer's ease of use with the unsubscribe request.
 - c. In order to comply with updated requirements for inbound direct mail to major mailbox providers, senders must implement one-click unsubscribe (List-Unsubscribe-Post) in the headers as described in [RFC 8058](#).

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 for an individual list. It is a best practice to default to removing the recipient from all mail.
8. 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.
9. Senders should use easily readable text descriptions instead of images to accompany hyperlinks to a one-click online unsubscribe webpage.
10. 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.
11. When a provider makes new subscription offers available, returning subscribers should be presented with the new selections unchecked 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 for recipients who use multiple email addresses forwarded to one central account or mailbox to prevent accidentally unsubscribing an address that doesn't belong to you. Laws surrounding handling of PII/personal data should be considered when implementing this.
14. The subscriber's email address should not be visible within the unsubscribe link itself. Links should utilize encoded references to protect personal data rather than displaying them in clear text.

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 [M³AAWG Vetting Best Common Practices \(BCP\)](#) for an in-depth guide to customer vetting.

6. 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

Infrastructure Configuration

- Tools that can be used to identify if a domain name has missing A or MX records: [What's My DNS?](#)

IP/Reputation Monitoring

- [DNS Twist](#) - Monitoring tool which can detect typosquatters, phishing attacks, fraud, and brand impersonation. Useful as an additional source of targeted threat intelligence and for defensive domain registration.
- [Google Postmaster Tools](#) - This is a free tool provided by Google to domain administrators who send emails to Gmail addresses. It provides valuable insights into email delivery performance, including delivery rates, spam rates, authentication failures (SPF, DKIM, DMARC), and more. Monitoring this tool helps identify and address issues affecting email deliverability to Gmail users.
- [Microsoft SNDS](#) - Smart Network Data Services (SNDS) The data on this site can help you build and maintain a good reputation for outlook.com accounts.
- [Sender Score](#) - Sender Score uses reputation-affecting complaints, traps, and deliverability data pulled directly from mailbox providers and message security agencies to generate your IP's Sender Score.
- [Talos Intelligence](#) - IP reputation and domain intelligence provided by Cisco Systems.
- [Yahoo Sender's Hub](#) - Yahoo's Sender's Hub provides similar functionality to Google's Postmaster Tools but is tailored for domains sending emails to Yahoo and AOL addresses. It offers insights into email delivery metrics and provides recommendations to improve deliverability. Monitoring Yahoo Sender's Hub allows administrators to optimize email-sending practices for better deliverability to Yahoo and AOL users.
- [Google Safe Browsing](#) - Safe Browsing is a service that Google's security team built to identify unsafe websites and notify users and website owners of potential harm.
- [Spamhaus DBL](#) - The Spamhaus Domain Blocklist (DBL) is a list of domain names with poor reputation. It is published in a domain DNSBL format. These domain reputations are calculated from many factors and maintained in a database, which in turn feeds the DBL zone itself.
- [SURBL](#) - SURBL specializes in reputation data provided in near real-time feeds. This data can be used to secure mail flow, mobile communications (SMS), web access (safe browsing, DNS Firewalls) and many other activities.

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, <https://www.iso.org/standard/75652.html>
- 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

- M³AAWG 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), <http://tools.ietf.org/html/rfc2369>
- RFC 3463: [Enhanced Mail System Status Codes](http://tools.ietf.org/html/rfc3463), <http://tools.ietf.org/html/rfc3463>
- RFC 5321: [Simple Mail Transfer Protocol](http://tools.ietf.org/html/rfc5321), <http://tools.ietf.org/html/rfc5321>
- RFC 5322: [Internet Message Format](https://datatracker.ietf.org/doc/html/rfc5322), <https://datatracker.ietf.org/doc/html/rfc5322>
- RFC 5598: [Internet Mail Architecture](https://tools.ietf.org/html/rfc5598), <https://tools.ietf.org/html/rfc5598>
- RFC 6376: [“DomainKeys Identified Mail \(DKIM\) Signatures”](https://tools.ietf.org/html/rfc6376) <https://tools.ietf.org/html/rfc6376>
- RFC 7208: [“Sender Policy Framework \(SPF\) for Authorizing Use of Domains in Email, Version 1”](https://tools.ietf.org/html/rfc7208) <https://tools.ietf.org/html/rfc7208>

Other M³AAWG Relevant Best Practices

- Messaging Anti-Abuse Working Group Position on Email Appending,
https://www.m3aawg.org/sites/default/files/doc_files/m3aawg_appending_position_update-2019-01.pdf
- M³AAWG Email Forwarding Best Practices,
<https://www.m3aawg.org/documents/en/m3aawg-email-forwarding-best-common-practices-version-2>
- MAAWG Vetting Best Common Practices (BCP),
https://www.m3aawg.org/sites/default/files/doc_files/MAAWG_Vetting_BCP_2011-11.pdf

Appendix B: Regulatory 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 ensure 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
- Europe:
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):
<https://ised-isde.canada.ca/site/canada-anti-spam-legislation/en/canadas-anti-spam-legislation>
- 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. (Compared to Transactional Messaging.)

Confirmation Message – (part of Confirmed opt-in process) An email sent to a new subscriber or recipient when they provide their email address to a sender. The recipient is required to click a link in the message to confirm they would like to be added to a list. If the recipient takes no action on the message they are not added to the list. Confirmation is usually obtained by clicking a link in the message or replying to it.

Confirmed opt-in (Double opt-in) - A process by which the sender confirms/verifies that the address submitted to them actually requested to be added to the list. For example, a user must click a link within a message they receive before they receive further communications (see Confirmation Message). Confirmed opt-in prevents users from being subscribed to a list without their consent.

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. (Compare to 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.

FQDN - Fully Qualified Domain Name is a domain name that specifies the exact location of network resources on the internet. FQDNs facilitate the identification and assist in access of online resources across networked environments. FQDN can contain multiple parts such as **www.subdomain.second-leveldomain.com**, where “.com” is the top-level domain and “www” is the hostname.

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 (MBP) – 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 the 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 provider's support or abuse desk, placing a phone call, or sending an email complaint to the provider's support or abuse desk, or to the sender of the message.

Opt-in – The process of a recipient indicating that they wish to receive 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 provider, or throttling of the connecting IP by the provider to slow down the rate of emails being delivered from that IP.

Subscription message /Welcome email - A message sent to a subscriber after signing up for an email list. The message usually contains information about the content and frequency of the mailings and should also include an unsubscribe link.

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. (Compared with Bulk/Marketing Messaging.)

Major edits for this version

- Major document restructure
- Sender ID deprecated
- Domain Keys (created by Yahoo) deprecated
- Identified Internet Mail (from Cisco) deprecated

As with all documents we publish, please check the M³AAWG website (www.m3aawg.org) for updates.