M^{3}AAWG | Messaging Malware Mobile
Anti-Abuse Working Group

Messaging, Malware and Mobile Anti-Abuse Working Group

M^{3}AAWG Help! I Hit a Spam Trap!

February 2023

The reference URL for this document is: https://www.m3aawg.org/help-i-hit-a-spam-trap

Introduction

This document helps Email Service Providers (ESPs) mitigate the consequences of hitting spam traps. It also suggests ways to use spam trap feedback to improve customers' sending practices, thereby minimizing future spam trap hits. In this document, "customer" refers to the organization using the ESP to send email.

Most email senders are faced at some point with the consequences of having sent mail to spam traps (or "spam trap hits"). The magnitude of the consequences can vary greatly depending on the number of trap hits, what type of trap was hit, who operates the trap, and other variables – all factors that customers may be unaware of. ESPs have a responsibility to monitor and inform their customers when a trap hit occurs. The ESP will want to prevent further trap hits and to mitigate the delivery effects of these hits. Failure to do so may lead to more severe consequences across the ESP's sending infrastructure.

A high rate of spam trap hits from a given mail stream can indicate an abusive sender, or at the least, a sender that may be inconsistent in their enforcement of best sending practices (M^{3}AAWG's recommended standards are laid out in M^{3}AAWG Sender Best Common Practices). The recipient domain may decide their users are best served by rejecting mail from that stream, or by assigning a lower priority to mail presented for delivery. In extreme circumstances, the ESP may find that they've been blocklisted and their mail rejected across a broad swath of the internet.

While hitting a spam trap is never desirable, spam trap hits can be used by the ESP as an opportunity not only to detect and remove abusive customers, but also to help their legitimate customers identify and correct poor sending practices. We cover mitigation and talking points for discussing spam traps with customers later in this document.

What Is a Spam Trap?

A spam trap is an email address used to collect, record, and monitor spam and other unsolicited or abusive email. Spam traps are designed to be indistinguishable from other email addresses and can be found across all types of networks, including corporate and "freemail" domains.

There are many different types of traps, but one thing they have in common is that they don't send email or subscribe to email distribution lists or newsletters. A trap operator monitors email sent to those addresses, and uses the data to analyze IP address and domain name reputation, as well as to evaluate email content.

These data are often used and redistributed via DNS blacklists (DNSBLs) and other reputation systems to help inform delivery or blocking decisions at recipient domains who utilize them.

A Taxonomy of Spam Traps

Some of the most common types of traps include:

Recycled Trap

An address or domain that may have been in active use by an individual at some point in the past, but that has since been retired after a period of inactivity and converted to a spam trap. The length of inactivity can vary significantly among operators, but M³AAWG suggests 12 months as a minimum. This sort of trap hit is usually an indicator of poor list management (or an ancient list), a lack of correct bounce processing, or both.

Pristine or Pure Trap

An address that has never been active for an email user before its deployment as a spam trap. A pristine trap hit is frequently the result of web harvesting, address space probing, or dictionary attacks. Pristine trap hits may be a strong indicator of the presence of purchased lists.

Typo Trap

This typically includes an intentional and perhaps common typographical error, most frequently in the domain portion of the trap address, e.g., user@gmial.com, user@notmail.com, and similar. This type of hit often indicates a failure of the customer to confirm the recipient's address, and can be caused by transcription or scanning errors committed at the time the sender collected the email address. While this may technically be a pristine trap, many operators classify these differently, as senders may obtain them through otherwise legitimate collection practices.

The M³AAWG document Best Current Practices for Building and Operating a Spam Trap provides a more detailed breakdown of the types of spam traps and their modes of use, and can be consulted for additional information. This document primarily references traps that directly influence delivery, as opposed to sensor trap networks used for reputation monitoring rather than mail blocking.

You know you've hit a spam trap when…

Since spam traps are commonly designed to be indistinguishable from other email addresses, it's often difficult to know when a given message has been sent to one. Indicators that spam traps are being hit can include the listing of a domain or IP on a blocklist or an increase in mail being rejected. Tools that monitor blocklists and reputation can provide metrics with regard to spam trap hits without disclosing (or "burning") the spam traps themselves.

On rare occasions, a recipient domain will announce the existence of a spam trap. This may be done via the host name of their MX server in DNS (e.g., spamtrap.domain.com) or text in an SMTP response stating the address is a spam trap.

There are also commercial spam trap feeds. These services, provided by deliverability monitoring companies, have their own spam trap networks. These networks are not used for blocking, but rather to allow customers of the deliverability monitoring companies to see how much of their mail is being sent to these traps.

Inadvertent Trap Exposure

Trap operators generally do not seek to expose their traps. The investment required to create and maintain traps that can produce useful data is significant. Operators must assume that once exposed, knowledge of a trap's existence will spread quickly, minimizing the trap's effectiveness.

In the course of investigation and remediation, ESPs or their representatives may inadvertently discover the identity of a spam trap IP address, domain, or network. To maintain the requisite stealthy operation, ESPs must take all appropriate steps to maintain the confidentiality of this data.

Customer communication should never explicitly or implicitly reveal the identity of traps or networks, and any discussion of such data by an ESP should be handled on a "need-to-know" basis.

In the event that identifying data about the trap or network is revealed and the identity of the trap operator is known, notifying the trap operator may be prudent. Informing the operator that their trap may be compromised will allow them to take any necessary action to maintain the efficacy of their network. Notifying the trap operator may also help an ESP establish or maintain a positive working relationship with the network owner.

Remediation

Customer Notification

An ESP has the responsibility to notify their customer when there is evidence a spam trap has been hit. Here are some considerations when notifying a customer of a hit.

Acquisition Audit

Spam trap delivery issues typically require an audit of the acquisition procedures that allowed the spam trap address to end up in the sender's database. Such an audit will necessarily resemble the list vetting procedures detailed in the M^{3}AAWG Vetting Best Common Practices document. However, some aspects of the audit will require a more granular approach, and will include further considerations such as:

- How were the contact lists created? An audit of the acquisition process should focus primarily on the methods used to acquire and verify the recipients within each contact list.
- Is it possible to learn when the sender first sent to a spam trap, and to correlate that event with a specific send or list segment to be targeted for review?
- Did the IP or domain appear on a blocklist as a result of the trap event, and is it possible to then infer what types of traps are implicated?
- Do some recipient domains appear with greater frequency in the implicated segment, indicating list poisoning or harvesting?

- Is the owner of the list willing and able to rebuild their list, and acquire permission from the recipients anew?
- Have previous sends resulted in previous block listings? If so, how were they resolved?
- Can the list owner identify the source of the problem data and remove all data acquired through that source?
- Can the ESP that was used to send the message obtain any additional data from the spam trap network owner?

As with initial customer vetting, key areas to examine in the course of an audit will include address collection, validation, and hygiene.

List Hygiene

As outlined in M^{3}AAWG's Sender Best Common Practices document, ESPs should review all feedback loop, bounce, and unsubscribe processing to ensure recipient addresses are being correctly processed and removed when needed.

ESPs should also review customers' list hygiene practices to mitigate the risk of sending mail to spam traps. Strict adherence to policies detailed in the Sender Best Common Practices document will result in the organic reduction of spam trap hits. Considerations should therefore include:

- Lower than average activity or engagement by recipients, within a particular domain, may indicate a spam trap network. If acquisition of addresses at that domain can be correlated with a particular list segment or acquisition method, then these list segments should be candidates for remediation, and the acquisition method should be discontinued.

- Senders should consider implementing a policy for suppressing recipients who remain chronically unengaged or non-deliverable. This minimizes the possibility of future trap hits should those addresses be converted into recycled spam traps. If the incidence of trap hits continue at their current level, the sender may consider adjusting their existing policy to make it somewhat more aggressive.

- Recent changes in list segmentation criteria or suppression file management may correlate with an increase in the rate of spam trap hits and should always be monitored closely. Special care should be taken when segmentation results in some recipients receiving mail who have not been sent to for an extended time, during which an address may have been retired and repurposed as a spam trap.

In any event, a customer or list that produces too many spam trap hits should be a candidate for thorough re-vetting. If the customer or list has already been subjected to a rigorous vetting process, then it's possible that a more recent change is complicit in heightened spam trap activity:

- Have there been any staffing changes at the customer's organization?
- Has there been a new API implementation, or recent changes in an existing API that may have created opportunities for API abuse?

- Have there been changes at the customer's address collection points that might indicate an abusable web form or opportunities for list poisoning?
- Have there been broader organizational changes like mergers and acquisitions activity or changes in business model that might indicate a need for a thorough re-vetting of the customer?

From time to time, an ESP may find itself with a client that refuses to participate in any remediation processes. It is strongly recommended that ESPs terminate customers who refuse remediation, and consider limiting access to data that might otherwise be provided to those customers.

Minimizing Future Incidents

Address Collection Practices

ESPs should review customers' means of address collection to identify any areas of concern. Collection practices sometimes incentivize consumers to share an email address without any form of data checking to ensure the address belongs to that consumer. These practices often lead to spam traps on lists. Higher risk collection practices include:

- Incentivized sign-ups
- Social media sign-ups
- Refer-a-friend forms
- Sweepstakes

Addresses collected using these strategies prioritize *any* email address over the *right* email address, and lead to poor quality lists.

Spam traps end up on lists in other ways as well, including:

- Typos from addresses entered at a point of sale
- Addresses harvested off websites (whether automated or manual)
- Purchased, rented or e-pended lists
- Lists of trade show participants.
- Single opt-in sign-up forms

During an investigation of address collection practices, ESPs should ask for specific opt-in data from the customer. A common investigative technique is to provide multiple addresses to the customer, including some addresses that are not on the list. The customer is asked to provide opt-in data, including:

- The time and date of the signup
- The URL of the form used and the connecting IP (if the signup happened online)
- The location of the transaction, if the address was collected in person.

Customers should be able to provide specific opt-in data, including but not limited to the URL for any web forms. As spam traps are sometimes added to lists by automated form submissions, customers and ESPs should review website traffic analytics. Non-traditional traffic or unusual spikes in volume may indicate a form targeted by bots, leading to an increase in spam traps on a list.

ESPs can also verify whether opt-in processes are functioning as intended. Does a sign-up at the URL, provided by the list owner at the time of the audit, result in a verifiable subscription? If it doesn't, consider the possibility that the URLs provided does not belong to the list owner, or that sign-ups from the page are shared among many. Verify whether any confirmation mechanism exists and works as intended.

Validation at Point of Collection

Customers can supplement good collection practices with email address validation strategies as part of their data collection process. Such validation may be undertaken via an internal review at the ESP, or by one of a number of services that provide validation on-demand or at the time of address collection. The best practice is to validate the email address as it is entered, and to prompt the subscriber to re-enter their address if validation fails.

When the validation is done in-house, there are a number of clues ESPs can look for when reviewing a list that might indicate poor or non-existent validation strategies. M^{3}AAWG Vetting Best Common Practices provides a comprehensive overview of these strategies.

Additional Prevention Strategies

Additional actions to consider to reduce future issues may include constraining the customer's ability to import lists, such as only allowing additions through an ESP-provided form script, or through some other process.

Sending restrictions such as sending only to engaged segments, or requiring the sender to delete or suppress segments that show little or no historical engagement, may also be applied as an additional step. Above all, the customer sender must be required to dispose of segments lacking confirmed permissions, although it may be acceptable to allow the customer to first attempt to reconfirm their permissions.

If a client is provisioned within a shared environment, it may become necessary to isolate that client on a dedicated infrastructure in order to minimize the potential for reputational damage to other senders using the same shared infrastructure.

Conclusion

When your system hits a spam trap, there are many things to consider. ESPs need to protect their infrastructure, but also need to assess customers who are hitting those traps. At the end of the day, the spam trap isn't the problem; it is a sign of an underlying problem with the customer's address acquisition and validation. Spam traps are a way to identify poor address collection techniques. The traps themselves are a marker that there are addresses on the list lacking permission. In many cases these addresses are going to belong to actual people who are now receiving spam. Fixing the collection processes that lead to spam traps also addresses the spam that is affecting real people. The main concern should always be with the recipients who may be getting spammed. There are many strategies that will work; this document provides a start at fixing the problem.

Reference

- On M^{3}AAWG's website, see generally "Documents for Senders and ESPs"
  https://www.m3aawg.org/documents-for-senders-and-esps
  and, in particular,
- M^{3}AAWG Sender Best Common Practices, version 3.0, updated February 2015
  https://www.m3aawg.org/sites/maawg/files/news/M3AAWG_Senders_BCP_Ver3-2015-02.pdf
- M^{3}AAWG Best Current Practices for Building and Operating a Spam Trap, version 1.2.0, updated August 2016
  https://www.m3aawg.org/documents/en/m3aawg-best-current-practices-for-building-and-operating-a-spamtrap-ver-120
- Vetting Best Common Practices (BCP), November 2011
  https://www.m3aawg.org/sites/default/files/document/MAAWG_Vetting_BCP_2011-11.pdf

As with all documents that we publish, please check the M^{3}AAWG website (www.m3aawg.org) for updates.