CLOUDMARK™

SMTP Abuse Prevention in IPv6 Networks

Positive Reputation Class of Service Method

June 2012

Michal Bujak

Foreword

This white paper presents a methodology for handling SMTP traffic from an unknown number of IPv6 source address senders with unknown intentions. Traditional blacklists and related technologies used successfully in IPv4 networks do not map well into IPv6 networks. This document describes the reasons why a substantially different approach to protocol-level abuse prevention in IPv6 networks will need to be developed.

An important aspect to consider with the recommendations presented in this paper is the fact that IPv6 SMTP transmission is essentially a clean slate. The difficulties in protocol-level protection in IPv6 networks can be overcome with the cooperation of SMTP originators. Effort must be made to strongly encourage major shifts to messaging server policy best practices that may have been impossible to effect in long established IPv4 networks.

Background

IPv4 address space exhaustion has caused a critical need for IPv6 deployment within fixed line operators, mobile operators, and enterprises alike. With a 128-bit addressable network space, IPv6 promises to support many more devices than would ever have been possible with IPv4.

Generally, service providers expect to allocate /56 to /64 prefix size ranges to each home or small business network customer entity (CE). Each CE would therefore have direct control over a network space of at least 64 bits within its assigned subnet. By comparison, the entire IPv4 Internet uses only 32 bits of address space.

It is clear that the size of IPv6 addressable space will prove problematic for some types of current anti-abuse solutions that worked very well in the IPv4 environment.

The Problem

The sheer size of the addressable IPv6 address space threatens to render useless many anti-abuse technologies that are based on IPv4 addresses. Many of the anti-abuse technologies used to protect SMTP servers in IPv4 deployments do not map well to IPv6 due to the enormous number of IPv6 addresses available:

- Long term IP reputation tracking
- IP blacklisting
- Traffic shaping

Many IPv4-based systems effectively track the reputation of IPv4 SMTP senders. These systems are capable of tracking the reputation of every address within the entire IPv4 space. Although problematic, identifying misbehaving senders and publishing corresponding blacklists is a tractable problem given today's computing resources.

However, the IPv6 space is many orders of magnitude larger than the IPv4 space. In the long run, it becomes practically impossible to track reputation across all potential sending addresses. As with IPv4, attackers will take advantage of this fact by quickly rotating through many individual IPv6 source addresses. Such high frequency address hopping allows the attacker to shed trivially any negative reputation associated with a particular IP address and quickly adopt a new address for which no such negative reputation yet exists.

To illustrate the potential scale of this problem, a single /64 customer network range could be used to launch a spam attack where just one message is ever sent from each unique IPv6 address within the range. Even with constant rotation of sending IP addresses, that attacker could theoretically sustain a rate of 580 billion messages per second without having to reuse an IP address in that /64 network more than once over the course of an entire year. Maintaining IPv6 blacklists in the same manner as IPv4, while potentially useful in the short term, will not scale in the long term as spammers retool and fixate on originating attacks from IPv6 connected sources.

Unlike blacklists, there are other IPv4 technologies that do not rely on the enumeration of IP addresses. Specifically, the Sender Policy Framework (SPF) and DomainKeys (DKIM) serve to authenticate the sending domain. These technologies enable tracking of domain-based reputation.

IPv6 Policy Challenges

In order to be effective, a defense strategy needs to be efficient and robust. Currently, it is estimated that 80-95% of all SMTP traffic over IPv4 is blocked at the connect stage with today's IPv4 blacklisting and throttling techniques. When an IP address, associated with bad traffic emanating from it, is identified, the connection can be closed immediately, avoiding the MTA overhead of further processing, such as reverse DNS lookups, cryptographic processing, directory or database queries, accepting the entire message for subsequent and much more expensive spam/abuse detection, authentication based on content (such as DomainKeys Identified Mail, or DKIM), and potential delivery into a spam folder or quarantine.

Ideally, a similar strategy would be followed for IPv6. However, given the scaling challenges of reputation tracking across all possible IPv6 addresses, we propose that the reputation of each sender be tracked primarily by its authenticated domain name rather than by its IP address.

Additionally, even though the potential threat space in IPv6 is much larger than that of IPv4, the number of legitimate senders in the IPv6 address space is not expected to be significantly larger than that of IPv4. As such, we will develop a methodology by which well-behaved senders are rewarded with greater MTA resources, rather than focusing solely on restricting poorly behaving senders.

In the discussion that follows, we define a sender identity as the identifier upon which the system can build reputation for senders. Two categories of identity will be discussed: domain-based identity and, if no domain-based identity is available, IP address-based identity taken from the connection context.

Below are presented the central principles of this white paper:

1. The system attempts to establish sender's identity.
2. All messages from unknown parties are forced, in aggregate, through an extremely limited default class of service (CoS). The default CoS constitutes a single shared pool of resources for which unknown senders must compete.
3. With proper identification, the MTA begins to track the sender's reputation. By identifying, the sender graduates from the default pooled class of service to a per-sender class of service that takes into account the sender's sending volumes and a measure of message quality.
4. Proper sender identification prevents a sender from stealing MTA resources from other senders by allocating a share of resources to each sender independently.

These principles constitute a significant departure from the traditional manner in which many MTA operators approach abuse prevention in their networks. In the sections that follow, we will express the level of access to MTA resources in terms of "class of service". Elevation or reduction of the CoS translates into greater or lesser access to resources, respectively.

The following graphic visually depicts the steps presented above:

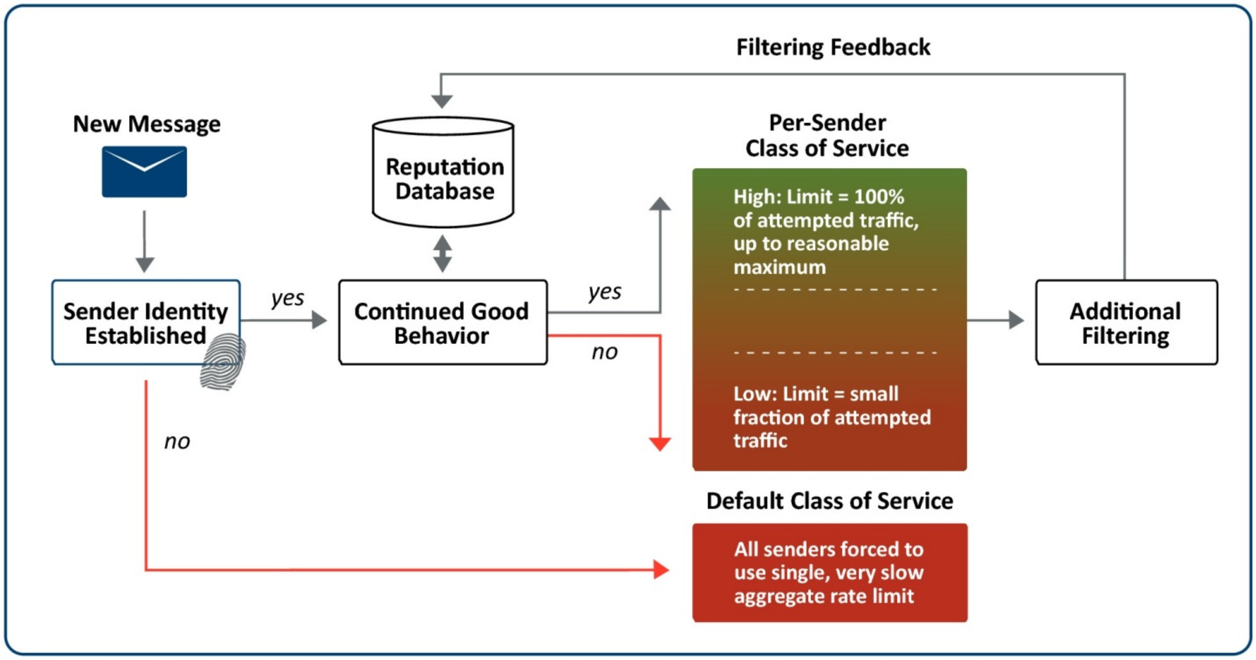
Filtering Feedback
New Message
Reputation Database
Per-Sender Class of Service
High: Limit = 100% of attempted traffic, up to reasonable maximum
Low: Limit = small fraction of attempted traffic
Sender Identity Established
yes
Continued Good Behavior
yes
no
Additional Filtering
no
Default Class of Service
All senders forced to use single, very slow aggregate rate limit

SMTP Traffic Shaping Through Class of Service

Class of Service Method for Handling IPv6 Senders

Historically, one of the most successful means of controlling the damage inflicted on a network by a sustained spam campaign is through traffic shaping. In the ideal case, trusted senders would be permitted a greater share of the MTA's capacity than senders with less trust. Senders identified as truly abusive would be given no capacity and blocked at the connection stage. Blacklists are used to identify poor and suspect senders.

In the IPv6 Internet, comprehensive blacklists do not exist as they do for IPv4. As such, we propose that with IPv6 SMTP messaging, all unknown senders be treated as suspect. By default, suspect senders are forced to share in a very limited pool of MTA resources in order to prevent potential attacks from doing too much damage downstream.

By identifying itself (through the use of DKIM[1] or SPF[2], WHOIS or WEIRDS[3] lookup, or known sender list), a sender will move up from the default class of service. The sender's reputation will be tracked and the sender will be assigned to a low per-sender class of service that may be elevated contingent on continued good sending behavior.

Each per-sender class of service level can adjust a variety of rate and throughput limits in order to control access to MTA resources (e.g., simultaneous connection limits, bandwidth limits, message rate limits, per-connection RCPT TO limits, etc.).

The following sections describe various class of service levels and the process by which the class of service is adjusted on a per-identity basis.

Low Class of Service - Default Throttling Policy

Until some form of reputation is computed for a sender, the SMTP client will be assigned to a combined default class of service applicable to all unknown senders. All SMTP clients that fall into the default CoS are forced to share in and compete for a single slice of MTA resources allocated to the default class of service rather than receiving any individual consideration. The explicit goal is to ensure that even if the vast majority of messages from unknown senders are spam, downstream damage would be kept to a minimum while giving additional filtering technologies time to react to an influx of spam messages.

The rate limits set in the default class of service may very well be lower than the corresponding limits in IPv4 SMTP transmission for similarly unknown senders. Such aggressive limits are necessary in order to encourage senders to take the necessary steps described in the next section.

Per-Domain Class of Service – Opportunities for Elevation

We propose that in order to graduate from the extremely restrictive default CoS, the system must be able to establish the sender's identity. Preferably, this would involve domain-based authentication, but other means of identification will be considered as well.

By identifying itself, the sender will have taken a minimum first step towards becoming a “good network citizen”. As a consequence of this effort, reputation for this sender identity will be tracked within the system. Additionally, the sender identity is rewarded with a per-sender class of service. The sender will no longer compete in the default resource pool, but will have its own class of service that may improve over time.

There are three general ways in which sender identity can be established. Domain-based identification is the preferred means of identification. However, points 2 and 3 below describe additional means by which senders who fail to sign may still be identified.

1. Using SPF and/or DKIM signing of mail
    - Results in domain-based identification
    - Domain-based identification is the preferred form of identification
2. Placement of the sending IP address or address range into a local “known sender list” either by manual means or automatically through additional reputation analysis outside the scope of this document.
    - Results in IP address-based identification.
    - Placement on this list automatically moves the sender into a per-sender class of service. This is not the same as whitelisting. Although the sender will be given greater access to the MTA's capacity, additional filtering technologies may still be applied.
    - Third party vendors, such as Cloudmark, are capable of providing a “known sender list” identifying known Email Service Providers (ESPs) and other, usually large, well-behaved senders.
3. A lookup in WHOIS or similar protocols, such as those being developed in the IETF’s WEIRDS working group, which can yield the IPv6 allocation to which the IP address belongs. If the allocation is too coarse-grained, its value for reputation tracking may be diminished because unrelated senders exhibiting different behaviors could be erroneously aggregated within a single measure of reputation. In this case, the information from WEIRDS could be disregarded.
    - Results in IP address-based identification.
    - Can result in the identification of a shorter IP prefix (i.e., a larger subnet) by which to identify the IP address range associated with a sender.
    - Additionally, the use of WHOIS may allow correlation of one IP block with another when it can be determined that they have a common registrant.

To a certain extent, identification constitutes a small hurdle that must be overcome. It is important to note that SPF and DKIM, in and of themselves, do not solve the abuse problem. In fact, spammers were early adopters of both technologies. Authenticated identities are used for reputation tracking and the subsequent elevation of class of service of good senders, but do not guarantee preferential treatment.

The use of SPF for identification is attractive because it can be invoked as early as the MAIL FROM stage. SPF provides the means by which the server can confirm that the client is authorized to use the MAIL FROM domain name. With proper SPF validation, the sender identifies its domain.

DKIM differs from SPF in that the entire message must first be accumulated prior to validation. In that respect, DKIM is much more expensive than SPF in terms of both network (the entire message must be accepted) and CPU resources (cryptographic checks are applied to the message). Regardless, with proper DKIM validation, the sender identity is based on the signing domain. DKIM is an important consideration because it can pass in cases where SPF does not, such as message forwarding.

Failures in SPF or DKIM can occur for valid reasons. Although it is tempting to degrade CoS or to drop the message outright, this can do more harm than good. Therefore, the current approach focuses only on successful validations.

With domain-based senders, the set of all IP address ranges associated with the authenticated domain would be maintained within the MTA's environment, and ideally, available in real time to all hosts in a cluster. The IP address ranges associated with a domain would be used to perform a connect-time check that skips the default CoS. However, if it is subsequently determined that the message is not signed, aggressive action against the message should be taken (e.g., drop, reject, etc.). If multiple domains use the same IP address, the message must be signed by one of the associated domains.

With address-based senders (i.e., domain-based identity is not available), all IP addresses that fall into a given IP address range will be associated with the given sender identity. The class of service for this sender would be applied across its constituent ranges.

Adjusting Class of Service - Traffic Shaping Policy

The operation of this traffic shaping approach is focused on rewarding good behavior while punishing irresponsible and outright bad behavior through incremental changes of per-sender class of service. An important reason for focusing on good senders rather than bad ones is that there are far fewer good actors than bad actors, thereby making the problem of identification and tracking much more tractable.

Upon graduation from the default CoS through identification, continued good behavior results in further elevation to higher classes of service on a per-sender basis. For the purposes of this discussion, we can define “good behavior” as the sustained sending of a high percentage of legitimate messages through the system as indicated by virus and spam filters applied to received messages.

A variety of approaches can be used for adjusting per-sender class of service. Here we present an abbreviated approach for illustration.

The rate limits imposed on a sender are raised in proportion to the expected traffic originating from a given sender.

Prior to accepting any part of a message, per-sender message limits will be applied at the connect stage. For an IP address-based sender identity, these limits are strict, corresponding exactly to the sender's allocated rate limits. For a domain-based sender identity, this is an aggregate limit, corresponding to the sum of rates associated with a particular IP address prefix.

The following steps summarize the administration of throttling policy:

1. The default rate for all senders in the default CoS is very low and shared among all such senders.
2. Senders can differentiate themselves through the use of one or more identification methods in order to graduate from the default class of service into a per-sender class of service.
3. The permissible rate, as with CoS elevation, would only be increased if most messages sent by the sender are classified as valid and not spam or virus.
4. In the case that message quality becomes skewed in the direction of spam, the permissible rate will fall along with class of service.

By following these steps, the system will be able to react dynamically to an influx of messages from known and unknown senders, rewarding good actors commensurate with their good behavior. It is important to remember that the vast majority of messages on the network are spam. This approach is still able to protect downstream entities in the absence of any information about a particular sender addresses.

Practical Considerations

In dual-homed IPv4 and IPv6 networks, the fallback behavior in case of IPv6 delivery failure is not clearly defined. When an SMTP client is aggressively rate limited in IPv6 and receives timeouts or outright message rejection, it could attempt to identify another IPv6 SMTP in an attempt to deliver again. The client could also fall back to IPv4 delivery. On the other hand, the client might do neither and simply retry later. It is important to keep in mind that, often, both IPv4 and IPv6 capabilities will be hosted on the same physical hardware. As a result, aggressive IPv6 throttling could precipitate additional MTA activity on neighboring IPv6 MTAs or IPv4 MTAs to which delivery could fallback.

Additionally, the methodology presented in this document can be enhanced in cases where an IPv4 fallback is available. If an MTA client is rejected by an IPv6-enabled MTA, and the client retries by sending the message to an IPv4-enabled MTA, all of the well-known IPv4-based abuse prevention techniques would then be applied.

As an optional optimization, a time windowed list of misbehaving IP addresses or IP address ranges can be maintained. This list can be used for blocking before the IP socket acceptance stage. If the MTA operates within a cluster, this temporal IP address blocking should be distributed across all hosts in the cluster in real time. The SMTP client should not be permitted to retry against a different IPv6 MTA.

Another issue is one of communicating errors related to class of service. As an example, if a connection is to be dropped, a 4xx response code can be returned, along with a text payload indicating the nature of the issue and instructions for remediation, such as a URI.

Summary

This white paper proposes an approach for fighting spam and other forms of abuse that does not lose efficacy within the much larger address space in the IPv6 messaging environment.

Rather than identifying an explicit list of addresses to block, this approach adopts the view that, by default, senders should receive a fairly low class of service, but are given ample opportunity to enhance their classes of service quickly by identifying themselves and exhibiting continued good behavior.

By requiring sender identification as a precondition to enhancement of class of service, the industry can encourage the adoption of best practices that have been difficult to effect in the long-established IPv4 world. In the best case, adoption of DKIM, WEIRDS, and heightened general awareness of the need to control spam, viruses, etc., could be driven by the recommendations in this white paper.

On a final note, it should be noted that nothing presented in this document explicitly ties the described methodology to IPv6. It would work equally well in the IPv4 space. The ideas presented take on greater relevance in IPv6 since the space is more unwieldy than the long understood IPv4 space.

References

[1] D. Crocker, T. Hansen, M. Kucherawy, "DomainKeys Identified Mail (DKIM) Signatures", RFC 6376, September 2011.

[2] M. Wong, W. Schlitt, "Sender Policy Framework (SPF)", RFC 4408, April 2006.

[3] Web Extensible Internet Registration Data Service (WEIRDS), https://datatracker.ietf.org/wg/weirds/