



BCP for Near Term SMTP Deployments in IPv6 Networks

August 2011

Kevin San Diego
Leon Rishniw
Murray Kucherawy

CONTENTS

EXECUTIVE SUMMARY	1
Background on IPv6	1
MESSAGING ANTI-ABUSE COMPLEXITIES INTRODUCED BY IPV6	2
Coexistence and Transition Technologies and their Impacts on Abuse	4
<i>Potential Issues with Dual-Stack Lite (DS Lite)</i>	4
<i>Potential Issues with Incremental CGN</i>	6
<i>Potential Issues with 6to4</i>	6
<i>Potential Issues with 6rd</i>	6
Domain Reputation Not a Direct Substitute for IP Reputation	6
ATTEMPTING TO ADAPT TODAY'S IPv4 MTA POLICY TO TOMORROW'S NETWORK	7
Adapting Allow Lists to IPv6	8
Adapting Block Lists to IPv6	9
Prudent Deployment of IPv6-Connected SMTP Servers	9
Track outbound Sender Reputation in Aggregate	10
Apply Class of Service Policy by IPv6 Prefix Length	11
Maintain IPv4-only Servers for uncontrolled Environments	12
INDUSTRY CALL TO ACTION	13
CONCLUSION	15
REFERENCES	17

EXECUTIVE SUMMARY

As unallocated IPv4 address space approaches exhaustion, IPv6 deployment becomes more critical for fixed line operators, mobile operators, and enterprises alike. Messaging industry understanding of how IPv6 impacts anti-abuse mechanisms is on the increase, but requires a broader acceptance of some of the risks involved with IPv6 connectivity and associated transition and translation technologies. While IPv6 solves the problem of limited IPv4 space, it also increases the chances of introducing security weaknesses within enforceable messaging server policy. IPv6 expands the threat surface for network attack in several dimensions and for messaging infrastructure IPv6 increases the threat potential by many orders of magnitude over the existing IPv4-enabled messaging infrastructure. This paper identifies the increased threat potential of IPv6, examines existing IPv4 based messaging security solutions, evaluates their applicability to IPv6, and provides recommendations for addressing the threats in an open, industry-supported manner.

This paper is intended for messaging professionals who are tasked with deploying IPv6 infrastructure within their networks. After reviewing this paper, readers should understand the likely threats that IPv6 deployment presents to their messaging infrastructure and be able to discuss rational deployment roadmaps with their management that support necessary messaging capabilities without exposing their messaging services to unnecessary risk and significant capital investment.

The messaging security community needs to establish new techniques and standards to support the evolution of messaging security in an IPv6-enabled world. Without industry leadership, ad-hoc solutions will arise that leave the network messaging infrastructure fragmented and vulnerable to abuse.

Background on IPv6

The IPv4 networks most likely to include PCs with bot infections are residential and small office/home office networks served by cable or DSL modems. These networks are typically assigned a single public IPv4 (32 bit) address and a compromised PC within this network is only able to send messages from that address. Anything sent from this PC can be isolated and tracked, triggering anti-abuse controls by a receiving messaging server. Even if the border network device reconnects to the service provider or refreshes its DHCP lease, in many cases this will not result in the network receiving a fresh IP address due to unexpired DHCP lease times. Contrasting this environment to IPv6 rollouts, most service providers expect to allocate large IPv6 network space (128 bit) to each customer entity, most commonly in the /56 to /64 prefix size range. If a PC is compromised within one of these networks, it could potentially source traffic from any IP address in that network's vast assigned IPv6 range. As IP-focused reputation systems are the norm in the IPv4 world, this high frequency IP address-hopping behavior can allow the bot to go undetected and shed any negative reputation it attracts after transition

to IPv6. 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 sent from each IPv6 address in the range. This means a spammer could send enormous amounts of spam without ever re-using a single IP address from inside its assigned network block. Even with constantly rotating sending IP addresses, that attacker could still theoretically be capable of sustaining 580 billion messages per second without having to send from any same IP address in that /64 network more than once over the course of an entire year. Moreover, re-training blacklists would be analogous to the per-IP blacklists we have with IPv4 now and has the potential to be a good start, but it still won't scale to IPv6 in the long run.

MESSAGING ANTI-ABUSE COMPLEXITIES INTRODUCED BY IPV6

Within today's widely deployed IPv4 messaging ecosystem, many of the effective and commonly commonly deployed anti-abuse mechanisms rely on evaluation of the IP address from which a connection attempt or message attempt is sourced. Some of these mechanisms include:

1. Reverse DNS verification (rDNS)
2. Real-time DNS Block List (DNSBL) checks
3. Bandwidth throttling based on Class of Service (CoS)
4. Sender Policy Framework (SPF) validation
5. Sender ID validation
6. Manual whitelisting
7. IP address warming
8. Other collaborative reputation systems

These general mechanisms involve comparing the origin IP address of a message delivery attempt to a policy (SPF, Sender ID) or a database of known good, suspect, and bad addresses.

Many of these anti-abuse techniques will be considerably crippled in a general migration to IPv6 due to the vast address space available. For example:

- Many service providers have declared an intention not to serve reverse DNS records for dynamic address space (even for small business networks), a requirement of MTA rDNS checks, given the memory and disk resources required to do so.

This will interfere with the myriad systems that leverage rDNS checks as a simple security feature, which includes ubiquitous Mail Transport Agents (MTAs) such as Sendmail.

- DNS servers providing DNSBL service will be required to accumulate and serve a set of data many orders of magnitude larger under IPv6 than is required under IPv4, with the obvious resource, caching, and performance penalties. Proposals for performing network aggregation within this data set, relieving some of this pressure, are few and nascent. It is possible the DNS will not scale even with the network aggregation schemes being considered, although they are certainly an improvement over the single IP address ideas.
- Today, traffic shaping solutions rely upon a sender having access to a relatively small number of possible IPv4 addresses and attempt to limit the damage possible from each IP address through bandwidth or connection shaping. These solutions become ineffective if attackers begin to spread message delivery load out over a vast set of source IP addresses.
- Data mining services that attempt to provide reputation services or IP address warming reports will begin to receive data about an enormous range of Simple Mail Transfer Protocol (SMTP) sources, most of which are malicious. Coalescing this information will be difficult as there are no defined standards for address allocation. The IETF attempted to standardize a typical IPv6 block allocation size, however it has recently retracted that position, stating that “the exact choice of how much address space to assign end sites is an issue for the operational community¹.” As there is no standardization or guidance on block allocation size, each service provider is free to subdivide its network allocation as it sees fit. Absent a standard means by which the service provider can securely indicate what allocation policies it implements, data mining services have no reliable information upon which to coalesce IPv6 addresses. This will result in reduced security due to conservative coalescing of IPv6 addresses or false positives due to aggressive coalescing of IPv6 addresses.
- SPF and Sender ID are in the best position to make a painless transition, though they are not on their own technically sound enough to obviate the need for continued reliance on other anti-abuse mechanisms or their future equivalents.
- Reputation systems that intend to rely on DomainKeys Identified Mail (DKIM) or other post-SMTP protocol message attributes are technically interesting, but in practical use provide none of the lightweight benefits of IP address based solutions currently deployed. For instance, for a receiving MTA to validate a DKIM signature,

¹ [Abstract, RFC 6177](#)

the message must first be fully accepted before the DKIM header can be authenticated, likely resulting in the squander of valuable bandwidth and MTA computing resources only to discover that the message is from a disreputable source based on a reputation lookup. There is no provision in the SMTP protocol to interrupt a message while it is being received by an MTA during the DATA phase. The only way to stop the sender from transmitting the remaining DATA portion is to drop the connection which causes the sending agent to assume it's encountering a transient problem and re-queue the message. Then the same message delivery attempt cycle begins again a few minutes later.

Coexistence and Transition Technologies and their Impacts on Abuse

Numerous technologies exist to aid the IPv4-to-IPv6 transition, including Large-Scale NAT (LSN, formerly known as "Carrier-Grade Nat"), Dual-Stack Lite (DS Lite), Incremental Carrier-Grade NAT (CGN), 6to4 gateways, 6rd gateways, 6in4 tunneling, Teredo relays, etc. Some of these technologies are in active use while others are experimental. Common themes among these include tunneling and address translation layers in various forms.

Some of these mechanisms obscure, to some degree, the actual IP address of the agent requesting delivery of a message. Others make it difficult to impossible for a receiving MTA to reliably understand how to track reputation for a sender traversing one of these layers. This obfuscation will prevent many present day IP address based security mechanisms from functioning effectively, as they are predicated on the IP address being a semi-static, unique identifier of a specific SMTP sender. In cases where the IP address in a packet has been replaced by that of a translation layer, tunnel or gateway, the ability to act on a specific source with accuracy is defeated. The following sections describe potential issues with three of these transition technologies.

Potential Issues with Dual-Stack Lite (DS Lite)

Dual-Stack Lite (DS Lite) enables service providers to deploy new customers with IPv6 only connectivity between the customer premise equipment (CPE) and the service provider's network. The appeal of DS Lite is that it supports customer environments that may still contain older computers and devices that do not support native IPv6 or dual stack addressing, while requiring the service provider to only provision an IPv6 address to the CPE layer. With DS Lite, an older IPv4-only device that is addressed locally via DHCP with an RFC 1918 address is able to connect to the Internet through a DS Lite-capable CPE device at the edge of the customer network. This router encapsulates the IPv4 packet in an IPv6 packet and forwards it across the IPv6 service provider internal network to an LSN layer. The LSN is then responsible for de-encapsulating the original IPv4 packet and NATing it out to the IPv4 Internet through a locally provisioned public IPv4 address. The resulting packet, now on the public Internet, has an IPv4 source address that belongs to the LSN. The packet is subsequently routed by downstream

IPv4 networks to the intended remote IPv4 service. In this case, all IPv4 SMTP traffic sourced from this DS Lite-enabled network will appear to arrive from a single shared external address bound to that service provider's LSN device.

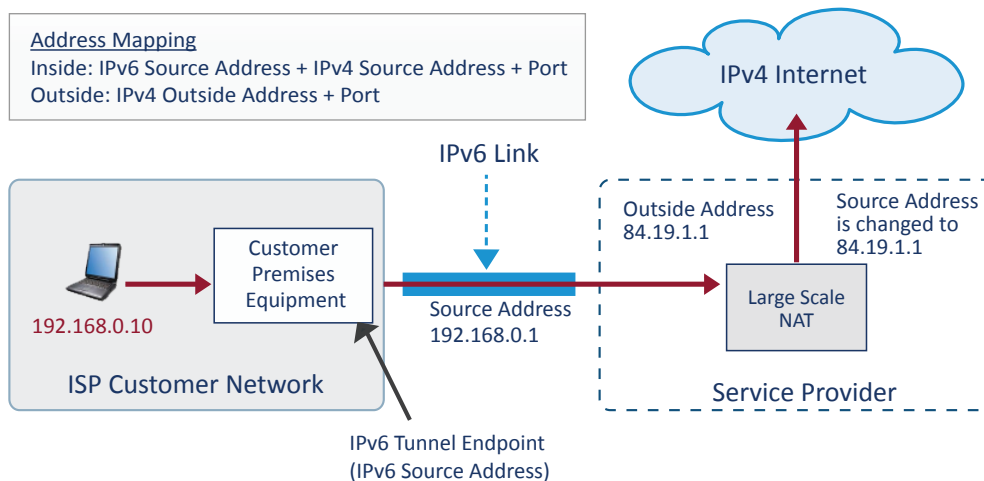


Figure 1
 Dual Stack Lite tunnels IPv4 packets over IPv6 between the user and the Large Scale NAT

An off-network SMTP server that receives such traffic will not be able to accurately ascertain the real source of the traffic, and will instead track reputation on the LSN's external IPv4 address. Because the remote SMTP server has no reliable mechanism to detect the IP address as a NAT device, it may mistakenly believe that the IP address represents a single IPv4 host. If sufficient bot activity exists within the service provider network that results in outbound spam traffic through the LSN, the LSN external IPv4 address has a high likelihood of being throttled or blocked by the remote MTA. One malicious bot behind traversing this LSN single IP address can harm the reputation of innocent customers who happen to route traffic using the same LSN external interface, resulting in service denial to particular remote services. This situation is already being experienced by mobile users whose traffic today traverses large capacity NAT devices, where one badly behaving mobile device can impact remote service accessibility for many other legitimate mobile users.

Moreover, should an abusive action warrant law enforcement attention, this obfuscation of the original IP address makes an investigation substantially more complicated as the various translation and tunneling layers must be unraveled to get at the true source of the abuse. This creates a burden on the agents operating these layers by requiring them to keep records about traffic transiting them; without such logging, tracing abuse back to its source becomes nearly

impossible. An initial set of best common practices² describing how service providers should be logging data traversing LSN layers is being discussed within the IETF presently though implementation will likely be fragmented.

Potential Issues with Incremental CGN

Incremental CGN is similar in function to DS Lite, with the exception that each CPE device communicates over the service provider network using IPv4 instead of IPv6. Any IPv6 traffic sourced from a computer system in a customer network is encapsulated within an IPv4 packet and forwarded to the LSN device for subsequent de-encapsulation prior to delivery to the final destination. For IPv4 traffic sourced from a computer system in a customer network, the traffic is directed towards the LSN device for NAT44 translation, effectively hiding the source address of the connection originator. The same issue with DS Lite of cross-NAT visibility for a remote SMTP server applies here as well, hampering the ability for a remote MTA to enact policy to effectively limit abuse from traffic originated from behind another operator's LSN device without causing collateral damage.

Potential Issues with 6to4

Considering another solution like 6to4, which enables any publicly routable dynamic or static IPv4 address on the Internet today to automatically claim a /48 IPv6 network range, the snowshoe³ problem of today's IPv4 messaging ecosystem can become a significantly larger problem for systems attempting to track reputation of individual IP addresses. Every 6to4-enabled network now has the potential for a bot to send large volumes of spam from many unique IP addresses within that /48.

Potential Issues with 6rd

While 6to4 gateways are easy to identify by their statically defined network prefix, 6rd (rapid deployment) gateways are not. Network operators providing 6rd service can assign any IPv6 network prefix to this gateway, preventing a remote MTA from easily identifying traffic sourced from that gateway as traversing a single 6rd gateway. There is also no mechanism for the target MTA to understand what the correct network prefix length is across which to aggregate reputation as this information is not published in a publicly accessible manner.

Domain Reputation Not a Direct Substitute for IP Reputation

Considering that most large messaging systems reject 80% to 95% of all IPv4 connection

² [draft-ietf-intarea-server-logging-recommendations-03](#)

³ Snowshoe spamming: The practice of a sender originating small numbers of spam messages across a large base of sending IP addresses in an effort to evade volume based detection by a receiving MTA.

attempts today due to IP reputation service listings or local behavioral history data, a connection level MTA policy provides significant protection against the most egregious, high volume senders at a relatively low resource cost. Domain reputation technologies that rely on authentication solutions like DKIM, where the digital signature required to authenticate the message to the purported sender is contained within the message header, require that each message is received by the MTA through the DATA phase of the SMTP protocol before the sender can be authenticated. Thereafter, additional analysis of the DKIM header needs to occur, including DNS lookup(s), cryptographic hashing of DKIM signature(s) against signed message header and body fields, etc. Carrying out this authentication process requires significantly more MTA and network resources to accept the connection and receive the message. Even if some of the typically heavyweight message acceptance functions (validating RCPT TOs against LDAP, running AS and AV engines, etc.) can be delayed until after the DKIM information is evaluated, this process would significantly increase work for the MTA rather than relying on IP reputation to make an early connection allow/deny decision. This results in increased infrastructure costs for the messaging platform.

Authentication of the sender verifies that the signing MTA associated with a domain handled the message and that the resulting message received by the local MTA was able to pass validation of included signatures. While domain reputation based on authenticated sending domains is useful data to track, relying solely on domain reputation on IPv6 enabled mail servers in lieu of IP address based reputation solutions will require more MTAs and network capacity to support the growth of abusive traffic (the anomalous last few months of 2010 excluded).

Domain reputation and DKIM are useful for abuse reporting, accounting of sending behavior of various legitimate marketing senders, anti-phishing (with ADSP and its successors), and handling of traffic from sources that require additional assurance of a message's source. To protect MTA systems from abusive senders at the connection level will require other solutions.

ATTEMPTING TO ADAPT TODAY'S IPV4 MTA POLICY TO TOMORROW'S NETWORK

In current IPv4 deployments, connection management based on IP reputation is the first line of defense and most resource effective method of protecting messaging infrastructure from the vast quantities of spam and other malicious traffic. This approach allows the operator to deploy a variety of CoS to their infrastructure, including denying connections, bandwidth shaping, connection concurrency limits, (re-)connection rate limiting, and RCPT TO limits. As messaging system operators consider opening their inbound environments to IPv6 SMTP connections, there are additional issues to consider due to the magnitude of increase of the network space. Some proposals recommend whitelists or "allow" lists where only approved senders can connect over IPv6 to the inbound MTA. Others believe continuation of single IP address blacklisting is sufficient to address messaging threats. The following sections describe these proposals.

Adapting Allow Lists to IPv6

First, consider a sender who is listed on the allow list. The simple policy would allow that IP address to connect to an inbound messaging server; however additional policies could be applied to the connection for CoS facilities. The first issue here is that IPv6 addresses have a local part, so there needs to be some level of aggregation below the full address. As a receiving MTA, where can it reliably look up the correct aggregation size for the connecting IP address? The natural inclination is to aggregate at the /64 level which represents the minimum subnet allocation. However this still poses significant technical implementation challenges in today's technology.

Secondly, consider a sender who is not listed on the allow list. There is no information on this sender and therefore the choice is simple: do not accept the connection. The expected behavior in this instance would be for the sender to roll through the various published MX records pointing to AAAA records, be blocked from connecting, and then fall back to connection attempts to MX records pointing to A records. This results in increased load on the receiver network and border MTA gateways to reject the same message delivery attempt multiple times. The fallback would be to attempt to deliver the messages over IPv4 connections. At this point we are back to the current state of messaging; having exhausted some amount of resources in applying the IPv6 policy. This presents some amount of resource depletion while accomplishing nothing as compared to an existing messaging infrastructure where default deny + allow lists are not in widespread use.

The derivative effect of a sender not on the allow list may also represent an increase in customer support issues. A sending MTA that is rejected as it cycles through published MX records may alert operators or trigger local sending policies that cause the messaging administrator to contact the receiving operator's support line to understand the reasoning behind the rejection, since the receiving IPv6 MX records are public information. The simple answer to this is to publish the information to the connecting MTA in the response code however this may only serve to magnify the support issue.

Allow lists also present a legal issue due to discriminating policies, such as that faced by Verizon in 2005 where the company was sued for blocking Europe, China and New Zealand from sending messages into their network. The result was significant customer backlash and a class action lawsuit. This precedent presents a challenge for any messaging administrator to push an allow policy through their legal department.

An additional consideration is the qualification of a sender for establishing sufficient credibility for inclusion on the allow list. Spamhaus⁴ currently provides such a whitelist that is restricted to IP addresses sending only transactional messages. The pre-requisite is that the IP address

⁴ [Spamhaus™Whitelist](#)

is not used for any other type of messaging, such as commercial transactional messages. As an operator looking to accept messages from the Internet at large, this presents a situation of discrimination against a new sender connecting to the network, where all new senders must essentially fall back to an IPv4 infrastructure. Such limited allow lists also do little to address the abuse vector presented by IPv6 as the number of senders who will not be on the allow list is large and they will have to fall back to IPv4 delivery after failing to connect over IPv6. In fact they would represent the majority of senders attempting to connect to the network once bot senders begin sending in larger volumes from IPv6 sources.

Adapting Block Lists to IPv6

Most blacklists that track spamming sources (excluding policy block list type services) need to track behavior for and generate lists containing millions to tens of millions of IPv4 addresses. In today's IPv4 messaging abuse landscape, it is relatively common to discover that a spammer has control of blocks of 256 IP addresses (8 bits) and greater. When IPv6 deployments expand in scope and size, the number of potential IP addresses that they can use to generate traffic will explode. For services that believe in tracking and listing bad senders at granularities where the majority of listings are a single IP address, this thinking does not map well to IPv6. A primary concern is the feasibility of collecting, calculating, and utilizing reputation information from IPv6 sources, many of which may never send more than one message, ever. A more sane approach would be to enable tracking of IPv6 spam sources of small networks by the network prefix size of that network. A core element of this type of functionality requires that a prefix size for a given remote network is publicly accessible in a system capable of high volume and low latency lookups. The only such proposal today seems to be from RIPE⁵. Network owners would need to voluntarily publish new WHOIS data elements corresponding to each network range's sub-allocation size. While this is a good first proposal, because it relies on network owners to properly maintain these records, the quality of the data will only be as good as the diligence and technical aptitude of the network owner. Additionally, because this information is being published to WHOIS it makes implementation of MTA policy challenging, as WHOIS systems are not built to support high volume querying and typically throttle queries by source IP. It should also be mentioned that Cloudmark is directly involved in efforts to add this capability at the IETF level, but the work is just getting started and the realization of these improvements is still years away.

Prudent Deployment of IPv6-Connected SMTP Servers

Introducing IPv6-enabled SMTP servers into large scale messaging environments has the potential to increase abuse by enabling attackers to source attacks from vast ranges of IP address space and exercise new methods for evading today's effective messaging server anti-abuse policy. Given the previously described methods for sender obfuscation via various IPv6

⁵ RIPE "Value of the "status:" and "assignment-size:" attributes in INET6NUM objects for sub-assigned PA space"

transition and translation technologies, it is critical for the continued security of messaging systems that initial acceptance of SMTP traffic over IPv6 networks be limited to environments where maximum sender behavior monitoring capability and MTA policy control are possible. Further, within these environments it must be possible for MTA policy to leverage consistent data to track behavior of senders over a period of time, enabling effective anti-abuse policy enforcement against malicious sources while minimizing collateral damage to legitimate senders.

In each of the following environments, the messaging system operator possesses the highest amount of control over the traceable data associated with each sender, namely an IPv6 source address contained within a larger allocated commercial or residential network of known prefix length and/or the authenticated user name provided during the initiation of an authenticated SMTP session.

- **Outbound SMTP servers accessible only to on-network, unauthenticated senders**
 - All connections are sourced from locally managed network entities (commercial or residential networks with known IPv6 prefix lengths)
 - MTA anti-abuse policy counters should be configured to track sender reputation by the appropriate IPv6 network associated with each network source, not by the individual source IP address.
- Outbound SMTP servers accessible only to on-network, SMTP Authenticated senders
 - All connections are sourced from locally managed network entities (commercial or residential networks with known IPv6 prefix lengths).
 - All SMTP sessions can be tracked by known authenticated user names.
 - MTA anti-abuse policy counters should be configured to track sender behavior by authenticated SMTP user ID (rolled up to the ISP parent account if possible) and by the appropriate IPv6 network from which the connection was initiated.

Track outbound Sender Reputation in Aggregate

Today's IPv4 residential networks are typically provisioned with a single dynamic or static IPv4 address. Many messaging systems in existing IPv4 environments today track sender behavior

by individual IP address or by authenticated SMTP user names, but new MTA policy capability is needed to track sender IPv6 reputation in a scalable fashion.

The number of potential source IP addresses associated with each residential network goes from manageable in today's IPv4 environment to intractable with IPv6 when IP-based sender reputation mechanisms are considered. With these now much larger customer network ranges and address autoconfiguration randomization methods such as Privacy Extensions (e.g. RFC 4941 vs. static address format outlined in RFC 2464), tracking reputation of abusive senders by /128 address becomes useless. Bot-infected systems within IPv6 residential networks are able to easily force the PC OS to cycle constantly through new IP addresses, sending only a single message from each IP address within a single residential /64 network, wasting resources on an MTA attempting to track reputation by individual IPv6 address.

The most effective method to handle this explosion of IPv6 addresses in customer networks is to treat each subnet as a single entity and track sending reputation by the overall IPv6 prefix of that subnet. Aggregating sending behavior in this manner enables reputation to be tracked and applied in a manageable fashion by the MTA. Connections originating from any IPv6 address within a given /56 to /64 IPv6 range will be associated with that source subnet's known IPv6 prefix. Any reputation checks and updates must be conducted against that overall network prefix based on observed sender behavior.

Solutions for the DNS based blacklist issue are nascent, but they do exist. One such proposal, which allows cache-efficient aggregation, is being circulated now within the IETF⁶.

Apply Class of Service Policy by IPv6 Prefix length

After compiling sending history of individual IPv6 network entities, MTAs should implement cluster-wide Class of Service (CoS) constraints for any associated IP addresses originating from a known IPv6 prefix. The MTA policy should optimally support assigning each IPv6 network entity to a given CoS based on observed behavior that matches a given CoS class entry criteria, and then apply the following constraints in a graduated fashion:

- Simultaneous connection limits
- SMTP protocol command "tarjit" wait periods (and ensuring PIPELINING support is disabled)
- Per-message RCPT TO limits
- Bandwidth limits

[6 draft-levine-iprangepub](#)

- Per-connection message limit (counted by mail data terminations, i.e. “<CRLF>.<CRLF>”)
- Per-connection SMTP protocol error limit (invalid commands, RCPT TOs, etc.)
- Messages per x seconds limits
- Recipients per x seconds limits
- Messages per y hours limits
- Recipients per y hours limits
- Connection retry mandatory waiting period

For IPv6 sending entities that have a history of sending a large percentage of good traffic, these senders would be assigned to a good sender CoS that allows higher limits. For senders who have been observed sending a mix of good and bad traffic, these senders would be assigned to a CoS that limits their sending behaviors significantly. For senders who have been observed to send almost no good traffic, these senders would be assigned to a CoS that nearly completely stops them from attempting message delivery.

For all other connections from senders with no previous history, these senders should be assigned to very strict CoS that limits message delivery until the MTA has an opportunity to learn a new sender’s behavior, as witnessed by subsequent traffic sourced from that IPv6 network prefix. MTA policy may even go as far as receiving several instances of content over a defined period and return a SMTP 4xx temporary failure for the sole purpose of scanning the content with an aim to classify the sender’s legitimacy without having to accept the actual messages.

Maintain IPv4-only Servers for uncontrolled Environments

In the immediate future, SMTP servers connected to the unrestricted Internet with only IPv4 connectivity will continue to prove easiest to defend. For IPv6 SMTP servers to be defensible in a similar manner to IPv4 SMTP servers, reputation technologies will require updates to various underlying Internet core systems to support precise application of reputation policy to connections from the wider Internet.

First, to accomplish reputation aggregation by tracking senders by their IPv6 prefixes, messaging systems must be able to query for the prefix length is for any given remote network from which a particular IP address is attempting to connect. At present, there is no reliable public mechanism to execute high volume, low latency lookups for a given IPv6 address’s prefix length to associate it with the network the address to which it belongs. Additionally, while many MTA systems today rely on reverse DNS to make policy judgments, many network operators

continue to struggle with a method to serve IPv6 PTR records in a scalable manner for their customer IPv6 allocations, rendering an MTAs attempts to continue to derive some data from DNS unreliable.

In lieu of a mechanism to track reputation precisely or to derive some information about a connecting IPv6 host via DNS, it is safest to configure all other SMTP systems with only IPv4 network interfaces to limit the attack surface and potential abuse:

- Inbound SMTP servers listed in external MX records, accessible from off-network SMTP sources despite the standard procedures established by several email RFCs including RFC974 and RFC5321, as this has been shown to be a reasonable anti-abuse measure.
- Outbound SMTP servers accessed by off-network, SMTP-authenticated users
- Outbound SMTP servers accessed by off-network Web Mail users (which are implicitly authenticated)

All delivery of inter-operator messaging traffic should traverse IPv4 network links, leveraging traditional legitimate sender announcement and authentication mechanisms.

INDUSTRY CALL TO ACTION

The potential threats created by IPv6 challenge the effectiveness of existing anti-messaging abuse techniques utilized with IPv4. More action is needed to establish and comply with a new set of best practices and techniques to mitigate these expanded attack vectors. Without industry involvement, it won't be feasible to limit or reject known bad sources early in the SMTP transaction, individually rate limit a given IP, or limit connections coming from a given IP. Solutions will likely involve standardizing or publishing CoS block allocations. IPv6 has been the subject of discussion in many industry forums from IETF (Internet Engineering Task Force) and APNIC (Asia Pacific Network Information Center) to NANOG (North American Network Operators Group) to MAAWG (Messaging Anti-Abuse Working Group). To date, the discussions of allocation and assignment have taken place during APNIC meetings however the discussions have been focused on assignment and allocation of the IPv6 address space. At NANOG, the focus has been on operational deployment issues. At MAAWG, an industry body focused on messaging security, there has been discussion on deploying IPv6 in the messaging infrastructure; however more focus must be placed on real threats to the messaging infrastructure presented by deploying IPv6.

Consider the impact on a high volume IPv4- based messaging environment where IP address level policies are removed. The connections to the network will heavily tax if not consume all network resources. All messages will need to be accepted before classification. The Class of Service attributes will only be applied after the message has been accepted. In short, the exist-

ing messaging infrastructure will be strained to the point of breaking, deliverability will suffer and customer satisfaction will deteriorate.

The advent of IPv6 in the messaging security community presents exactly this challenge. Reputation systems for IPv6 must be elevated to an aggregated level. This requires industry agreement to drive standardization and adherence to Class of Service block allocations or at a minimum, comprehensive and reliable publishing of aggregation block sizes per IPv6 range. This will facilitate reputation at the CIDR level which can be applied at connection time and therefore not require message acceptance through part of a message delivery transaction.

In Europe, for example, RIPE is allocating a minimum /32 IPv6 allocation⁷ for local Internet registries (LIR) with a minimum /64 sub-allocation⁸ for individual entities requiring a single subnet and any entity requesting more than a /48 requiring justification for the request. There has been little consideration for abuse particularly in the messaging sector where snowshoe attacks are commonplace from an IPv4 /24 CIDR subnet (or larger). The potential of this to explode beyond control when snowshoe attacks originate from an entire IPv4 space (IPv6 /96 CIDR) needs to be considered and addressed at the industry level.

While IPv6 has been discussed in industry forums such as APNIC, NONOG, and MAAWG, none of these discussions have focused on the implications of IPv6 for messaging abuse. Without industry collaboration on developing and applying best practices, consistent, effective policies will be impossible.

Currently, there are no DNSBLs for IPv6 that contain any useful information. DNSBLs form a major part of the security policy in IPv4 mail infrastructure. Without aggregation, the management of a comprehensive list of IPv6 addresses and reputations becomes untenable due to the vast increase in data. Secondly, consider a sparse DNSBL and the resulting default policy for a connecting IPv6 address. The most aggressive CoS would allow one message to be sent over the connection. With the size of standard allocation blocks it will be feasible to open numerous connections to the infrastructure from within a single allocation and still get through any CoS policies.

Without some form of publicly available, reliable, aggregation standard there is no reasonable way to apply connection limiting for a specific CIDR block as there is no way to determine if an IPv6 address belongs to a larger allocation than a /64. There have been discussions on assigning reputation by nearest neighbor, where the reputation of an unknown connecting IP address is associated with its nearest known neighbor. This policy is basically guesswork and will result in inappropriate reputation and the resulting service impact will increase support costs. Existing systems should be extended to provide this information such that connecting IPv6 addresses can be quantized to their appropriate allocation range where reputation can be applied.

As operators rush to deploy IPv6 there is a natural pressure to support all services on the IPv6 infrastructure. The current trend is to deploy dual-stack implementations to endpoints in the network to ensure IPv4/6 interoperability. This means that all endpoints will have IPv4 capability for the foreseeable future. Rushing deployment of messaging services on IPv6 without

⁷ [RIPE minimum /32 IPv6 allocation recommendation for LIRs](#)

⁸ [RIPE recommendation on individual network entities](#)

adequately planning and testing messaging security policies exposes the carrier network to unnecessary risk. Rather than focusing on operational deployment issues, the industry urgently needs to identify potential threats to messaging and carefully deploy IPv6 based technology on that knowledge. The longer the industry waits to define standards, the more challenging the solutions will be.

CONCLUSION

This paper has discussed the need for IPv6 migration in messaging, the challenges with existing security policies, and the limitations on these policies implied by current migration technologies. We've provided recommendations for initial rollout of messaging in an IPv6 environment and for messaging industry alignment in an IPv6 world. A high level threat analysis has been presented along with a recommendation for industry alignment on best practices for messaging in an IPv6 world.

As the need for deploying IPv6 becomes mandatory, there is significant pressure to migrate all services over to IPv6. Messaging professionals must consider the security implications of IPv6 to develop rational IPv6 migration plans that don't compromise their internal messaging systems. Messaging professionals need to drive the IPv6 migration by establishing a roll-out plan for their upgrades that provides the necessary functionality without compromising their internal messaging systems. All messaging rollout plans should include an analysis of the threat surface along with technology selection. A mapping of existing messaging security policies onto an IPv6 infrastructure must be an early step in the planning process so that the viability of each policy can be assessed in the context of IPv6.

However, introducing messaging services to external IPv6 networks in messaging infrastructures will require further Internet evolution; if we are to enable a similar level of protection to what currently exists in IPv4 networks. This includes the ability to ascertain a remote network's default IPv6 prefix allocation size.

Messaging professionals must consider the security implications to IPv6 to develop rational migration plans that do not compromise internal messaging systems. Additionally, they must be capable of sizing the impact of any migration plan on their infrastructure. This should involve a thorough evaluation of the security systems available with IPv4 to determine how the same services can be provided after the transition to IPv6. For elements that cannot be afforded similar protection under IPv6, IPv4 should remain the preferred interface until the environment has sufficiently evolved. This is a crucial and important first step in the planning process.

This paper is a call to action for the messaging security industry. Without early discussion, consensus on best practices and adoption of the recommended best practices, operators worldwide will adopt individual policies that inhibit messaging or compromise messaging infrastructure security. While many industry bodies have held IPv6 sessions over the past years, there needs to be a shift in focus from education to pragmatic security policy discussion. MAAWG

appears to be best positioned to make messaging security recommendations in conjunction with ICANN.

Deploying IPv6 messaging infrastructure will be required of every operator in the coming years. With the current level of expertise operating IPv6 based messaging infrastructure, a major challenge is accurately assessing the security risk implications of any IPv6 deployment decision. This becomes increasingly complex in an environment where the primary method of security analysis is based on extrapolating experience with IPv4 messaging threats onto IPv6 infrastructure. As such, there needs to be a realistic appreciation of the limited view of the threat landscape posed by IPv6 in messaging that is used by each operator when assessing their IPv6 deployment plans.

REFERENCES

1. Postel, J., Internet Protocol, [RFC 791](#), September 1981
2. Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. J., Lear, E., Address Allocation for Private Internets, [RFC 1918](#), February 1996
3. Crawford, M., Transmission of IPv6 Packets over Ethernet Networks, [RFC 2464](#), December 1998
4. Carpenter, B., Moore, K., Connection of IPv6 Domains via IPv4 Clouds, [RFC 3056](#), February 2001
5. Nordmark, E., Gilligan, R., Basic Transition Mechanisms for IPv6 Hosts and Routers, [RFC 4212](#), October 2005
6. Hinden, R., Deering, S., IP Version 6 Addressing Architecture, [RFC 4291](#), February 2006
7. Huitema, C., Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs), [RFC 4380](#), February 2006
8. Lyon, J., Wong, M., Sender ID: Authenticating E-Mail, [RFC 4406](#), April 2006
9. Wong, M., Schlitt, W., Sender Policy Framework (SPF) for Authorizing Use of Domains in E-Mail, Version 1, [RFC 4408](#), April 2006
10. Allman, E., Callas, J., Delany, M., Libbey, M., Fenton, J., Thomas, M., DomainKeys Identified Mail (DKIM) Signatures, [RFC 4871](#), May 2007
11. Narten, T., Draves, R., Krishnan, S., Privacy Extensions for Stateless Address Autoconfiguration in IPv6, [RFC 4941](#), September 2007
12. Siemborski, R., Melnikov, A., SMTP Service Extension for Authentication, [RFC 4954](#), July 2007
13. Klensin, J., Simple Mail Transfer Protocol, [RFC 5321](#), October 2008.
14. Despres, R., IPv6 Rapid Deployment on IPv4 Infrastructures (6rd), [RFC 5569](#), January 2010
15. Allman, E., Allman, J., Delany, M., Levine, J., DomainKeys Identified Mail (DKIM) Author Domain Signing Practices (ADSP), [RFC 5617](#), August 2009 Narten, T., Huston, G., Rob-

- erts, L., IPv6 Address Assignment to End Sites, [RFC 6177](#), March 2011
16. Durand, A., Droms, R., Woodyatt, J., Lee, Y., Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion, [draft-ietf-softwire-dual-stack-lite-07](#), March 2011
 17. Jiang, S., Guo, D., Carpenter, B., An Incremental Carrier-Grade NAT (CGN) for IPv6 Transition, [draft-ietf-v6ops-incremental-cgn-03.txt](#), January 2011
 18. Levine, J., An efficient method to publish ranges of IP addresses in the DNS, [draft-levine-iprangepub-01](#), December 2010
 19. Perreault, S., Yamagata, I., Miyakawa, S., Nakagawa, A., Ashida, H., Common requirements for IP address sharing schemes, [draft-ietf-behave-lsn-requirements-01](#), March 2011
 20. St Sauver, J., MAAWG IPv6 Session, [PDF](#), February 2009
 21. Hogewoning, M., Van Mook, R., Value of the “status:” and “assignment-size:” attributes in INET6NUM objects for sub-assigned PA space, [ripe-513](#), February 2011
 22. Durand, A., Gashinsky, I., Lee, D., Sheppard, S., Logging recommendations for Internet facing servers, [draft-ietf-intarea-server-logging-recommendations-03](#), February 2011
 23. Doyle, J., Understanding Carrier Grade NAT, [Network World Blog](#), September 2009
 24. Doyle, J., Understanding Dual-Stack Lite, [Network World Blog](#), October 2009

For more information
visit us at www.cloudmark.com

Americas Headquarters
Cloudmark, Inc.
San Francisco, USA

Asia Pacific Headquarters
Cloudmark, Inc.
Singapore

Europe Headquarters
Cloudmark Europe Ltd.
London, UK