IPv6 Prefix Discovery

Murray S. Kucherawy
<msk@cloudmark.com>

Today

- IPv4-based email abuse prevention relies heavily on a database of IP addresses with "bad reputations"
- A database of addresses is no larger than 4.3 billion entries (of course)
- Most popular expression of these is the RBL (Realtime Block List, RFCxxxx), which is published via the DNS

RBLs

- Query: 1.2.3.4.rbl-root, ask for "A"
- Reply: NXDOMAIN, or 127.0.0.1
  - Or sometimes the octets in the reply encode reputation data
- Caching and redundancy keep this functional and practical
  - …so far

IPv6

- Vastly larger address space
- Not practical to consider tracking reputation about each of them
- No standard delegation size; commonly between /48 and /64
- A spammer could send junk from such a network and rarely, if ever, re-use a single address

RBLs under IPv6

- Ignoring database size for a moment, this still won't work
- A spammer changing IP address quickly will mean caching of previous answers becomes useless
- And other cached data will be flushed because of space limits
- So this would clobber the DNS in general

What's needed

- We need to be able to figure out , given an IP address, the size of the endpoint delegation
- Allows address aggregation by reputation systems
- Keeps the query space about the same as it is for IPv4 now
- IRTF has an idea out there that allows the DNS to express IP ranges

Some ideas

- Publish it via WHOIS
  - WHOIS isn't standard and doesn't seem scalable
  - Some registrars can't be trusted to publish real data
- Get it from BGP
  - MTAs don't really have access to BGP data
  - We'd need a standard interface to exchange it between the lower layers and the higher ones

Can you help?

- Does this working group's mandate fit the idea of exploring this?
- Or does it belong in some other WG?
- How would you suggest we go about doing this?