Using BGP for realtime import and export of spam whitelist/blacklist entries

Peter Hessler
phessler@hostserver.de

Hostserver GmbH

17 September, 2015

traditional methods

network-based spam fighting:

- download a file from a server every $periodic
- live lookups from an external provider (e.g. DNS lookups)

traditional methods

some obvious problems with both methods

- only as fresh when you downloaded the file
- ...the provider may only generate the file on their own schedule
- ...leading to most-pessimistic schedules
- massive load and congestion at the top of the hour
- “network bogons” problem
- ability to receive mail is limited by the external service response speed
- behavior when the service is not available (e.g. Spamhaus ddos)

bgp-spamd.net

network-based spam fighting:

- bypass and trap lists from spamd(8)
- use BGP-4 and BGP communities (RFC 4271 & RFC 1997) for distribution and labeling

bgp-spamd.net

- only list the specific IP addresses that exhibited a specific behaviour
- do *NOT* penalize/reward network neighbors
- really simplistic, we just want to catch the low-hanging-fruit
- don't open your mail server to the world
- don't block the world from seeing your mail server
- greylisting is powerful, when it still applies!

why is this useful

- use the bypass and trap lists from 3rd parties
- ...they are much larger than you (and/or)
- ...they have different traffic patterns than you
- ...semi-trusted servers are usually semi-trusted elsewhere
- ...ditto for attackers
- shared bypass lists help the “gmail sender” problem

bgp-spamd.net

- available at `http://www.bgp-spamd.net`
- all configurations and scripts are available
- I am interested in additional “spamd-source” servers, please contact me
- and of course, more users are always welcome

bgp-spamd.net

- Publically launched at AsiaBSDCon 2013 on March 17
- 3 upstream sources
- 4 users

bgp-spamd.net

- A year later (16 May 2014)
- 5 upstream sources
- 28 users

bgp-spamd.net

- 6 months ago (14 March 2015)
- 5 upstream sources
- 55 users
- 2 route servers

bgp-spamd.net

- Today (12 September 2015)
- 5 upstream sources
- 134 users
- 2 route servers

spamd-source trap list

- using greylisting
- generated from source server's spamd trap list
- addresses are listed if their first delivery attempt is to a spamtrap
- expires in 24 hours from last delivery attempt

spamd-source bypass list

- spamd has a very low bar to be added to the whitelist
- ...redelivery within 4 hours
- ...kept in the whitelist for 36 days.
- semi-trusted email server list used to bypass spamd
- higher entry bar than normal spamd whitelist
- in the whitelist for 75 days, and sent more than 10 emails
- ...we “think” it’s a real mail server
- again, do not be overly aggressive

results

SUCCESS

lessons learned

- overall, a success
- generally positive reactions from users

the good

- many sources sharing information
- block lists are supurb

the good

- 3rd parties are making this work with non-OpenBSD users!
- Mark Martinec made it work with FreeBSD, rbldnsd, and SpamAssassin
- Anonymous using Quagga and their Propritary infrastructure
- (thank you!)

the good

- very fast to update
- 7 seconds to download the full bypass and trap lists over crappy home dsl
- 2 seconds to propagate changes to all members
- ... can be even faster, needs more work

the bad

- bypass list has too many spammers on it
- ... several users have mentioned they had to stop using it
- ... we need to spend more time adjusting the heuristics

the bad

- server crash, causing 5 day outage
- ...while I was on vacation (in New Zealand)
- ...and during long holiday weekend in the US
- ...where the only route server was

the ugly

- I have not been as responsive as I should have been
- have not had a lot of time to dedicate to improving
- ... code
- ... sources
- ... client usage

future work

- fix the heuristics for addition to the bypass list
- ... a bit *too* relaxed
- the “gmail sender” problem is *worse* with IPv6!
- ... a single email can use many hundreds of IPs within the same /64

future work - spamd-source

- easier processing of spamd(8) on spamd-source systems
- can spamd differentiate how it received the data
- more spamd-sources from different and new countries
- ... University students in CA do not send a lot of email to JP

future work - brainstorming

- voting
- ... “two upstreams think an IP is X, then make it X”
- ... somewhat tricky, as BGP doesn’t support this
- (ab)using an RPKI lookup process to process addresses before adding
- ... pretend to do RPKI
- ... do stuff
- ... allow or disallow address from being listed
- ... for now, only thoughts with no code

Acknowledgements

Many thanks to
my coauthor Bob Beck,

- the University of Alberta at `ualberta.ca`
- Bob Beck of `obtuse.com`,
- Henning Brauer of `bsws.de`
- Peter N.M. Hansteen of `BSDly.net`,

for being sources of `spamdb` information.

- `Sonic`

for hosting the California USA implementation `us.bgp-spamd.net` and

- `Hostserver GmbH`

for sponsoring the Frankfurt Germany implementation `eu.bgp-spamd.net`

Questions?