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

a year in the life

Peter Hessler
phessler@openbsd.org

OpenBSD

16 May, 2014

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

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

bgp-spamd.net

- Today (16 May 2014)
- 5 upstream sources
- 28 users

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

spamd-source

- 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!

spamd-source trap list

- 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

why is this useful

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

statistics - also known as 'blatent lies'

- 163,608,694 Events

statistics - also known as 'blatent lies'

- for now, only IPv4 entries
- none! of the IPs can be aggregrated into non /32 netblocks
- 17,088 /24's are represented in the bypass list
- 22,275 /24's are represented in the traplist

traplist statistics - also known as 'blatent lies'

- 1,174,783 unique addresses
- 31 entries with more than 10k announcements
- 10,346 entries with more than 1k announcements
- 801,001 entries announced only once

traplist statistics - also known as 'blatent lies'

Unique Unroutable IP Addresses

- 3 entries from 0.0.0.0/8 ('this' network)
- 20 entries from 10.0.0.0/8 (RFC 1918)
- 1 entry from CGN Shared network
- 2 entries from localhost (127.0.0.0/8)
- 1 entry from 169.254.0.0/16 (link local)
- 7 entries from 172.16.0.0/12 (RFC 1918)
- 1 entry is 192.168.0.0/16 (RFC 1918)
- 15 entries are "Multicast" (224.0.0.0/4)
- 17 entries are "reserved" (240.0.0.0/4)
- total of 9849 additions

traplist statistics - also known as 'blatent lies'

Top 10

- 1 13,958 78.83.35.125/32 78-83-35-125.spectrumnet.bg.
- 2 13,735 39.53.191.193/32 ptcl.net.pk.
- 3 13,735 36.76.49.143/32 telin.co.id.
- 4 13,718 46.161.117.1/32 adsl-46-161-117001.crnagora.net.
- 5 13,625 79.29.46.11/32 host11-46-static.29-79-b.business.telecomitalia.it.
- 6 13,594 5.22.65.170/32 asiatech.ir.
- 7 13,573 84.54.160.128/32 bginfo.net.
- 8 13,559 2.182.8.116/32 dci.co.ir.
- 9 13,559 2.180.24.90/32 tci.ir.
- 10 13,486 81.169.146.190/32 mo4-p07-ob.smtp.rzone.de.

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

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

the ugly

- still no IPv6 support
- ... well, the distribution mechanism works perfectly fine
- ... “just” needs spamd(8) support

lessons learned

- overall, a success
- generally positive reactions from users

future work

- fix the heuristics for addition to the bypass list
- ... a bit *too* relaxed
- (still) add IPv6 support to spamd
- 36 hour days

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 - route-servers

- second route server in Europe
- more reliability over lossy networks

future work - brainstorming

- voting
- ... “two upstreams think an IP is X, then make it X”
- ... somewhat tricky, as BGP doesn’t support this
- deeper level of integration between bgpd and spamd
- ... partial syncs of spamd databases
- ... spamd use pf tables for all the things?
- ... for now, only thoughts with both upsides and downsides

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.net`

for hosting the reference implementation `rs.bgp-spamd.net`

Questions?