Spam blocking
methods and experiences

Linuxdays Luxembourg 2003

christian mock

<christian.mock@vibe.at>

http://www.tahina.priv.at/~cm/talks/spamblocking.{sxi,pdf}

version 1.3

contents

- how spam is sent
- different methods for blocking spam
  - reject during SMTP dialogue
  - RBLs and local blacklists
  - SpamAssassin
- experiences with various techniques
  - percentages of hits on my mail system
  - methods to use and methods to avoid

first things first: risk analysis

- a.k.a. "know that you will be losing mail, and why"
- know your user's requirements
- test (e.g. by tagging) before you block
- monitor effectiveness
- block on hard criteria, tag on fuzzy
- do you want to block spam or LART luser admins?

how spam is sent

- direct (spammer -> recipient MX): ISPs will mostly terminate spammer's account immediately
- that means spammers need to hide their tracks to keep their accounts
- forged headers intended to cause confusion
- sender address mostly fake or stolen
  - that means filtering on sender address makes little sense

open relays, open proxies

- send mails to some unrelated server, let that one do the work and it's admins handle the trouble
- lately, open proxy abuse is on the increase
- HTTP proxies support "CONNECT" (to tunnel SSL connections)
- CONNECT smtpserver:25 HTTP/1.0
- leaves no trace of spammer's IP address in mail headers

and it gets even more stealthy

- viruses/worms and IE exploits (e.g. in spam "unsubscribe" pages) install backdoor on broadband-connected PCs
- spammers use those "zombies" to
  - send spam
  - DDoS anti-spam sites
  - run nameservers + web-redirectors for the spamvertized sites
  - the involved zombies change every 5 minutes

rejecting during the SMTP dialogue

- (all percentages are % of rejected RCPTs, Oct 2003)
- technical criteria: HELO
  - sender must give HELO (0%)
  - check HELO parameter syntax (2%)
  - don't accept HELO with own hostname/IP address (7.7%)
  - don't accept "localhost"/"localhost.localdomain" (0.6%)

rejecting during the SMTP dialogue

- SMTP pipelining: only when negotiated (0%, used to be more)
- sender domain must exist (7.6%)
- check header/body for asian charset declarations (0.6%)
- don't accept for unknown local users (7.8%)
  - catchall domains are dangerous: dictionary attacks
- don't relay (0.4%)

local blacklists

- can be based on sender domain, client hostname's domain, client IP address
- block countries by IP space (extreme measure)
  - china (4.6%)
  - korea (3.5%)
  - taiwan (1.1%)
  - hongkong (0.5%)
- block some ISPs by client host's domain (5.8%, 6 ISPs)

RBLs ("Real-Time Blackhole Lists")

- work via DNS, e.g. for 209.88.103.4:
- 4.103.88.209.proxies.relays.monkeys.com
  IN A 127.0.0.2
  IN TXT "BLOCKED: See
  http://www.monkeys.com/upl/
  listed-ip-0.cgi?ip=209.88.103.4"

RBL types

- based on different criteria:
  - open relays: relays.ordb.org, relays.visi.com
  - open proxies: opm.blitzed.org
  - fed from spamtraps, by country, operator's preferences, ...
- quality assessment may be difficult
- you depend on an EXTERNAL source
  - osirusoft RBL closed down due to DDoS and blacklisted *all IPs*

SPEWS

- taking attitude re-adjustment to a new level
- anonymous, communication via NANAE newsgroup
- lists IP ranges of known spammers
- "intentional collateral damage": expands listings (shortens netmasks) if ISP doesn't react
- listed ISP's users are supposed to pressure ISP to kick spammers

RBLs at work

- list.dsbl.org (19.4%)
- cbl.abuseat.org (11.5%)
- SPEWS (7.9%)
- opm.blitzed.org (7.4%)
- relays.visi.com (5.0%)
- sbl.spamhaus.org (4.4%)
- blackholes.easynet.nl (2.5%)
- relays.ordb.org (0.4%)

more RBL info

- List of Lists:
  - http://www.declude.com/junkmail/support/ip4r.htm
- quantitative comparison:
  - http://www.sdsc.edu/~jeff/spam/Blacklists_Compared.html
- intro to blacklists
  - http://www.scconsult.com/bill/dnsblhelp.html

RBL tools

- online-checkers for lotsa RBLs:
  - http://rbls.org/ http://openrbl.org/
- build-your-own tool
  - http://spfilter.openrbl.org/

content analysis

- best way to detect spam, IMHO
- mail must be received in full
- can check different properties
- based on a combination of properties, better decisions on spamminess are possible

bayesian filters

- gets trained on samples of spam and non-spam
- computes probability of single words in spam/non-spam
- checks mail and calculates "spamminess" probability based on words in mail
- needs continuous training on user-specific material, but is very effective

bayesian filters

- idea and first paper by Paul Graham
  - http://www.paulgraham.com/spam.html
- standalone: bogofilter
  - http://sourceforge.net/projects/bogofilter
- SpamAssassin >2.50
- ASSP Anti-Spam-SMTP-Proxy
  - http://assp.sourceforge.net/
- and more...

Razor, pyzor, DCC

- principle: users report spam to a database, others query that DB
- "fuzzy checksumming" methods run over mail body, checksum is reported and queried
- razor2 implements "reputation scheme" for spam reporters

SpamAssassin

- http://www.spamassassin.org/
- perl, open source, Unix/Windows
- gives a score (positive/negative) per property
- sum of scores > threshold: spam detected
- more than 800 tests
- very configurable and extendable
- supports Razor, Razor2, DCC, pyzor

SpamAssassin

- SpamAssassin checks:
  - header inconsistencies
  - Received: header lookup in RBLs
  - characters sets used
  - language (heuristic detection)
  - text fragments
  - MIME structure (syntax, HTML without text/plain, ...)

what to do with detected spam?

- /dev/null ??
  - nobody can notice false positives
- tag, and store into "junk" folder ??
  - who's got the time to regularily read it?
  - mail gets lost anyways
- generate bounce ??
  - with all the faked sender addresses...
- reject during SMTP
  - spam gets dropped, but sender will notice on "honest" false positives

my setup

- mail server based on postfix serving about 20 users and 10 mailing lists
- running a combination of all techniques mentioned
- few false positives
  - monitoring still needed
  - whitelisting is also essential
- about 1-5 spams in my inbox daily, ATM

negative experiences

- some RBLs
  - ORBS did arbitrarily list people they didn't like
  - dial-up RBLs give lots of false positives
  - beware of RBLs closing down on short notice!
- clueless postmasters
  - a customer's mail-partner was listed as open relay, long fixed, but never bothered remove

negative experiences

- filtering for hostname without domain in HELO
- filtering on client IP without reverse DNS
- filtering Received: headers against RBLs and dropping mail
- setups must be adapted to mail system user's requirements

statistics

Period: Oct 5 - Nov 4, 2003

Non-Spam mails: 2578

Rejects vs non-Spam:

RBLs, SPEWS	249%		
SpamAssassin	142%	RCPT checks	33%
HELO checks	44%	Sender Domain	32%
Country BLs	41%	ISP BLs	25%

done.

thanks for your patience

questions?

<christian.mock@vibe.at>

http://www.tahina.priv.at/~cm/talks/spamblocking.{sxi,pdf}