

SPAMIKAZE

Devdas Bhagat
devdas@dvb.homelinux.org

What is Spamikaze?

- Spamikaze is a spamtrap driven DNSBL generating system.
- A spamtrap is an address or a domain which receives only spam.
- Spamikaze builds the list of spam source **IP addresses from mail headers.**

Motivation

- Need for a DNSBL with low false positives.
- Need for a DNSBL with low maintainance.
- Avoid the need for secret spam recognising techniques.
- Should be able to work with any level of spam even if log files are not available.

Features

- **Low false positive rate** due to spamtrap driven nature.
- **Auto expiry** of addresses.
- **No questions asked** removal technique.
- **Whitelisting** of SMTP clients is available.
- Allows for **backup MX** servers.

Core technology

- Just a bunch of GPLed **Perl** scripts.
- **Database driven backend** – currently supports MySQL and PostgreSQL.
- Perl CGI for removing addresses and reporting.
- Currently depends on **parsing the contents of a MH style mailbox or maildir** to generate the IP address list via a cron driven script.

Why the PostgreSQL port?

- The MySQL code stores IP addresses as raw integers in four columns – bad for performance.
- **MySQL's limitations** rendered the code **incapable** of scaling up to large loads.
- A dual Athlon with 2 GB of RAM was not able to handle a feed of a few dozen IP addresses.

PostgreSQL port benefits

- Major code and database structure reorganisation.
- **Fewer queries** needed to complete the transaction.
- Native IP address type used.
- Can be used as a **DNSBL generating system**, or as a **direct IP address map** for those MTAs which support it.

Current status

- Basic setup working.
- **DNSWL functionality** is in CVS.
- Reorganisation of MySQL codebase to reflect the Pg codebase.
- **Daemonised version of the parser** for live inserts (single threaded version done, multiplexing version based on Net::Server in progress).

Vapourware

- Multisite sharing of DNSBL data sources.
- Multisite sharing protocol uses RFC 2822 message bodies, with transport left to NNTP and SMTP as appropriate.
- Alternative methods of detecting legitimate mail sending hosts.

Multisite requirements

- **Distributed trust** design.
- Ability to have **different versions** supported concurrently.
- Ability to **resist various DoS attacks** (both content based and packet flooding).

Links

- <http://spamikaze.nl.linux.org/> -- Home site and MySQL codebase
- <http://nixcartel.org/~devdas/spamikaze-pg.tar.gz> -- Pg codebase (being merged into CVS)
- <http://psbl.surriel.org/> -- Spamikaze driven public DNSBL.
- spamikaze@nl.linux.org -- Mailing list.

Questions and Feature Requests

?