Hermes Anti-Spam Proxy

The Hermes SMTP Proxy and Implemented Anti-Spam Techniques

Author: Veit Wahlich

EMail: veit AT ruhr.pm.org

Date: 09.06.2007

http://hermes-project.org/

Hermes Anti-Spam Proxy

What is Hermes?

- transparent SMTP proxy
  - to be placed in front of real SMTP server
  - accepts connection from the Internet, relays to real server
  - monitors communication between client and server
  - minimal impact on communication
    - only generates rejects and disables PIPELINING extension
    - intercepts STARTTLS and SMTP authentication extensions
    - inserts processing information into email header
    - should not interfere with any other SMTP communication

Hermes Anti-Spam Proxy

What is Hermes?

- implements a wide range of anti-spam techniques
  - active checks
    - greylisting
    - DNS-based blackhole list (realtime blackhole list)
    - reverse DNS check
  - passive checks
    - banner delay
    - throttling

Hermes Anti-Spam Proxy

What is Hermes?

- active SMTP extensions support
  - supports STARTTLS connection encryption extension
    - required to monitor encrypted communication
  - supports SMTP authentication extension
    - disables active checks such as greylisting on authenticated connections

Hermes Anti-Spam Proxy

What is Hermes?

- more facts on Hermes
  - coded in C++
  - uses SQLite 3 as database backend
  - developed by ITEISA, S.C. (El Astillero, Spain)
  - first public release in April 2007
  - source available under terms of the GPL
  - packages for many platforms
    - generic binary RPMs for i386 Linux systems
    - generic source RPMs for easy recompilation (RPM-based Linux distributions, Apple Darwin, etc)
    - binaries for MICROS~1 Windows

Anti-Spam Techniques in Hermes

- assumptions Hermes works on:
  1. spammers have to deliver as many emails as fast as possible to work profitable
  2. the number of simultaneous connections per host is limited
  3. the number of hosts is limited
  4. it is impossible for spammers to track transmission status data for every recipient

Anti-Spam Techniques in Hermes

- SMTP banner delay
  - passive check
  - the initial SMTP banner (code 220) is being delayed for a configured time
  - protocol enforcement
    - the protocol requires the connecting SMTP client to wait for the banner before sending any data
    - any protocol compliant SMTP client will wait (legitimate servers)
    - spammers under time pressure will try to send data without receiving the 220 banner or simply drop the connection
  - clients not respecting the protocol get a 20 seconds penalty before Hermes closes the connection
    - extra annoyance for spammers

Anti-Spam Techniques in Hermes

- throttling
  - passive check
  - a 1 second delay is inserted after every SMTP command
    - throttles down spammers and blocks their connections
    - many spammers drop connection if delivery takes too long
  - throttling stops as SMTP server accepts the DATA command
    - email content is transmitted at normal speed

Anti-Spam Techniques in Hermes

- greylisting
  - active check
  - integral and most complex function of Hermes
  - combination of blacklisting and whitelisting
    - soft reject and blacklist for a short time on first attempt
    - accept and whitelist for a longer period on second attempt
    - first email is received with a delay of few minutes, follow-ups arrive without delay
  - rejected clients get a 20 seconds penalty before reject is sent
    - extra annoyance for spammers

Anti-Spam Techniques in Hermes

- greylisting
  - Hermes implements strict greylisting
    - tracks composition of sender and recipient email addresses and client IP address
    - only exact matches get whitelisted
    - causes problems with Google Mail (work-around exists)
  - static whitelisting capabilities exist
    - exclude single recipient addresses or even whole domains from greylisting
    - exclude hosts or whote subdomains from greylisting based on reverse DNS hostname

Anti-Spam Techniques in Hermes

- DNS-based blackhole lists
  - active check
  - rejects mail if client's IP address is listed in configured DNSBL
    - either automatic (realtime) lists generated from registries (i.e. dial-in networks), open relay scans, caught by honey pots etc.
    - or handpicked lists of known spammers
    - i.e. Spamcop.net, NJABL, Spamhaus, ORDB and blackholes.us
  - effective – but error prone
    - innocient ISPs, mail and web servers get blacklisted constantly
    - as of version 1.3, Hermes does not support more than one RBL
  - rejected clients get a 20 seconds penalty before reject is sent
    - extra annoyance for spammers

Anti-Spam Techniques in Hermes

- reverse DNS check
  - active check
  - rejects mail if client's IP address does not resolve reverse
  - widely-used and -accepted technique – but susceptible
    - soft-rejects mail if all your name servers fail
    - no standard or RFC requires clients to offer reverse resolution
  - rejected clients get a 20 seconds penalty before reject is sent
    - extra annoyance for spammers
  - not included in vanilla Hermes 1.3 source release
    - will be included in future releases
    - available as patch from mailing list archives
    - already included in current RPM builds and SRPM

Enhancing Hermes

- Nolisting with Fakehermes
  - fake SMTP servers always dismiss mail with soft rejects
  - protocol-aware SMTP clients will try delivery to every server referenced by an MX record with ascending distance
  - many spammers either spam only the first MX record (lowest distance) or the last one (highest distance, “backup MX”)
    - backup MX servers often hold less extensive anti-spam facilities
  - even expandable to an “Unlisting array”
    - stands for analysis of sequence order
    - allows auto-blacklisting of hosts acting incompliantly
    - allows auto-whitelisting in greylisting table for incoming mail
    - currently not directly supported by Fakehermes

Enhancing Hermes

- content-based filtering
  - adding third party software such as SpamAssassin or DSPAM to your MTA will rate spammy mails that pass through Hermes
  - may add estimation filters on resp. evaluation of
    - rating contained words and de-obfuscation
    - layout of and tags used in HTML emails
    - IP addresses in DNSBLs and parts of hostnames from headers
    - sender domain verification (i.e. SPF, DomainKeys, SenderID)
    - sender email addresses
    - header abnormalities
    - statistical distribution analysis (i.e. Bayesian estimation)
    - invoking checksum clearinghouses (i.e. Razor, Pyzor, DCC)
    - and much more...

Hermes Anti-Spam Proxy

Any questions?

Hermes Anti-Spam Proxy

Thanks for your attention!

Resources

- Hermes and Fakehermes
  - http://hermes-project.com/
- PDF of these presentation sheets and patches (including Veit Wahlich's reverse DNS check patch)
  - http://ruhr.pm.org/material.psp
- Apache SpamAssassin
  - http://spamassassin.apache.org/
- DSPAM
  - http://dspam.nuclearelephant.com/