Hit the Ground Spam(fight)ing

LISA ’05, San Diego

December, 2005

John “Rowan” Littell   Earlham College

littejo (at) earlham (dot) edu

$ARGV[0]

- There is no magic bullet.

- Many products, both commercial and open source; the best ones combine methods and have feedbacks among methods.

- Time is short – I won’t mention everything.

- Goal: Help you fight spam or understand the systems that are doing it for you.

Two Approaches

Protocol Hacks

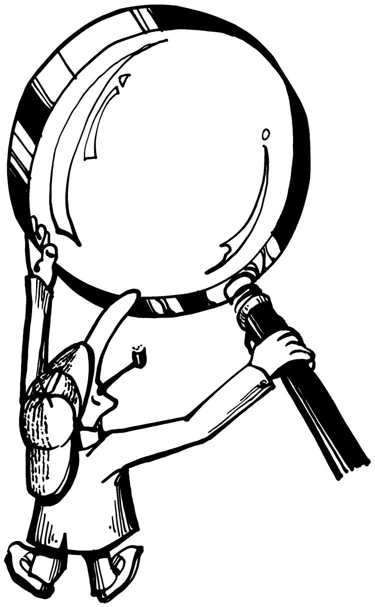

Content Analysis

Protocol Hacks: DNSBL

- Reject mail from IP addresses presumed to be spammers via DNS lookup
  - Pros: Quick, widely supported
  - Cons: Quality varies, false positives can be hard to work around
  - Suggestions: Choose a well-respected one, have a method in place for exceptions
  - Reference: http://en.wikipedia.org/wiki/DNSBL

Protocol Hacks: Greylist

- Tempfail the first instance of sender/recipient/IP address triplet, accept when it tries back
  - Pros: Entirely within SMTP, effective against virii
  - Cons: Delays the first message, some broken SMTP servers don't play well
  - Suggestions: Choose a flexible one, use the well-known whitelist, have a method for exceptions
  - Reference: http://en.wikipedia.org/wiki/Greylist

Protocol Hacks: Callback

- Test that sender address can receive mail via MX probe
  - Pros: Basic form of address verification, mostly hidden from users
  - Cons: Can create backscatter with MXs that blindly accept any address
  - Suggestions: Set exceptions for servers that don't work well with callback, have a method for other exceptions
  - Reference: http://www.snertsoft.com/sendmail/milter-sender/

Protocol Hacks: TMDA

- Require senders to validate themselves the first time they send to a recipient
  - Pros: Very effective
  - Cons: Outside of SMTP protocol, requires a fair amount of human maintenance, some implementations may be defeatable by automatic methods
  - Suggestions: Pre-whitelist your regular correspondents and any approved non-human senders or use pre-certified tagged addresses
  - Reference: http://en.wikipedia.org/wiki/TMDA

Protocol Hacks: SPF

- "Reverse MX" – check that mail originates from valid IP for sender domain
  - Pros: Mainly a method for tracing the sender of a message (no guarantee about nature of content)
  - Cons: Breaks basic forwarding and "aliases" lists, only useful in most restricted cases, slow adoption
  - Suggestions: Reject or raise the score on messages that fail to comply with published SPF records
  - Reference: http://en.wikipedia.org/wiki/Sender_Policy_Framework

Protocol Hacks: SMTP and TCP Tricks

- Require senders to follow RFCs and basic good behavior
  - Possible Methods: HELO before data, HELO string checking, Sendmail “greet pause”, throttle connections, reduce bandwidth
  - Pros: Catches a number of spamware systems
  - Cons: Catches a few legitimate mail server implementations, some methods need maintenance, some methods are only implemented as hacks (milters, etc.).
  - Suggestions: Watch for exceptions, don't use high-maintenance “tricks”

Content Analysis: Accept/Reject Lists

- Sender addresses (or patterns) to accept or reject
  - Pros: Regex patterns can match a number of spamware senders
  - Cons: Requires maintenance
  - Suggestions: Comb logs for identified spam and add regex patterns, keep regexes simple, allow users to build their own basic lists, use auto-whitelisting

Content Analysis: Content Matching

- Simple (keyword) or complex (regex) matching of spam content
  - Pros: Can be very effective (e.g., SpamAssassin), keywords are easy for users to understand
  - Cons: Regex rules are complicated and need constant tuning or updates, only catches known spam content
  - Suggestions: If using keywords, allow users to specify them themselves; couple content rules with other methods or use a scoring technique

Content Analysis: Fuzzy Signatures

- Compute fuzzy checksums of messages to compare with known spam content
  - Pros: Can be fairly accurate, can work around minor obfuscation techniques
  - Cons: Only able to recognize known spam content, relies on others' identification of spam
  - Suggestions: Use as a scoring technique, heavy users of free services should join submission network
  - Implementations: Distributed Checksum Clearinghouse (DCC), Vipul's Razor

Content Analysis: SURBL

- Spam URI Realtime Blocklist – DNS based URI list
  - Pros: Effective against phishing or other click-through spam
  - Cons: Only works with known spam URIs
  - Suggestions: Use as a scoring technique and give matches a high weight
  - Reference: http://www.surbl.org/

Content Analysis: Bayesian Classification + Learning

- Calculate probability of spam content based on learned spam words and tokens
  - Pros: Over time can become very accurate, requires little maintenance
  - Cons: Diverse mail content can lower accuracy
  - Suggestions: Allow users to build individual Bayes databases for individual accuracy, combine with site-wide database for shared known spam
  - Reference: http://en.wikipedia.org/wiki/Naive_Bayes_classifier

Content Analysis: Antivirus

- Identify known e-mail viruses and executable content
  - Pros: AV engines are very accurate for viruses, some include phishing matching
  - Cons: Takes resources
  - Suggestions: Dump or quarantine positive matches, **do not** send sender notifications – **this is spam!**

Where to Can Your Spam

- **Client or Access Server** – perform content analysis in the MUA or the POP/IMAP daemon
  - Pros: getting to be easy for users
  - Cons: can only do content analysis
  - Suggestions: use as a first step or if your e-mail provider won't support other methods
  - Examples: POPFile, Thunderbird, MacOS X Mail

Where to Can Your Spam

- **Mail Server** – integrate spam processing as part of incoming or final delivery on primary mail server
  - Pros: easy to set up, easy to tune for individual users
  - Cons: heavy load on server, final delivery not the best place for processing
  - Suggestions: best suited for small sites
  - Examples: SpamAssassin called by procmail, numerous milters and plugins for Postfix, Qmail, etc.

Where to Can Your Spam

- **Spam Gateway Appliance** – insert appliance as primary MX
  - Pros: easy to install, automatic updates and shared signatures, best place for spam processing, quarantine areas, redundant units
  - Cons: cost and quality varies, some work best only with certain mail architectures
  - Suggestions: verify recipient addresses, restrict access to internal mail servers
  - Examples: Postini, Barracuda, Mirapoint Razorgate, CanIT, Meridius, build-your-own. . .

Where to Can Your Spam

- **Firewall/IPS** – inline security appliances that can perform content analysis or protocol hacks
  - Pros: can protect entire network, some can operate as an invisible bridge
  - Cons: new technology, header tagging and quarantine often not available
  - Examples: OpenBSD's spamd(8)

Eating Your Spam

- The Recipient Is Correct

  . . . except when the President, the Lawyers, or Your Boss is. . .

  COROLLARY:

- There is an Exception to (nearly) Everything

  . . . be prepared, technically and politically, to deal with exceptions for filter hurdles. . .

  HOWEVER:

- The Recipient Needs Simplicity

  . . . don't give your average user too many configuration options. . .