What spam?
(with postfix and spambayes)

Derrick Hudson

February 23, 2004

Contents

1 Background

At the time of this writing, my home server is the only one I manage. It receives on the order of 700 messages per day with about 160 rejected by SMTP–level restrictions. Another 68 are discarded by the server due to identification of the content as being some form of Microsoft Windows executable. Finally, 30–50 more junk messages are identified and discarded during delivery through the use of a Bayesian classifier. All the remaining spam (very little, about 3–5 messages per day) is delivered to my "junk/unsure" folder. A few non-spam (about the same amount) is also delivered to that folder from imperfections in the Bayesian classification.

As you can see, this system is very effective and quite pleasant because I actually see almost no spam whatsoever even though my email address is readily available on the web in several forums and archives. The effectiveness derives from layering multiple defenses against junk mail. These layers are composed of postfix[1], maildrop[2], and spambayes[3].

2 postfix

To begin, note that Jim Seymour maintains a very good explanation of postfix' anti-UCE controls at [5]. I highly recommend reading his paper too.

2.1 "boilerplate"

```
smtpd_delay = yes     # default
```

2.2 smtpd_recipient_restrictions

Recipient restrictions are applied at the "RCPT TO" stage of an SMTP transaction. This is the last point at which the server can prevent unauthorized relaying. Therefore, be very careful about when and what to 'OK' in the following tests.

2.2.1 postfix:

```
smtpd_recipient_restrictions =
        reject_unauth_pipelining
        reject_non_fqdn_sender
        reject_non_fqdn_recipient
        reject_unknown_recipient_domain

        check_sender_access hash:/etc/postfix/access-sender
        reject_unknown_sender_domain
        check_recipient_access hash:/etc/postfix/access-rcpt
        check_recipient_access pcre:/etc/postfix/access-rcpt.pcre
        permit_sasl_authenticated
        # permit_tls_clientcerts
        check_client_access hash:/etc/postfix/access-client
        permit_mynetworks
            # do the pcre test first in order to declare RFC
            # non-compliance before other HELO reasons
        check_helo_access pcre:/etc/postfix/access-helo.pcre
        check_helo_access hash:/etc/postfix/access-helo
        reject_invalid_hostname
        reject_non_fqdn_hostname
            # /after/ the HELO is sanity-checked,
            # see if aol.com, etc., senders are using a legit
            # sending server
        check_sender_access hash:/etc/postfix/access-freemail-sender

        reject_unauth_destination
        check_recipient_maps

smtpd_restriction_classes = check_freemail_access
check_freemail_access =
             check_client_access hash:/etc/postfix/access-freemail-client
             reject
```

2.2.2 English:

Reject pipelining on non-ESMTP connections (only abused proxies or netcat(1) will trigger this)

Reject (envelope) sender addresses that don't have a fully-qualified domain name (without a fqdn no mail can be sent back, so the "address" isn't worth receiving from)

Reject (envelope) recipients that lack a fully-qualified domain or the domain is unknown. No legitimate mail will have an incomplete or erroneous recipient. If the sender made a typo, then they are informed of the problem sooner rather than later (the message can't be delivered anyways).

Reject sender domains that are not known. This is a DNS check to see if the domain the sender claims to be in exists. If it doesn't, then the address is invalid and not accepted.

Check the recipient address against an explicit blacklist table (one is a keyed lookup, the other a perl-compatible regular expression match).

Permit authenticated senders. (This is ASMTP) If a sender has authenticated, and the sender and recipient address pass the above sanity checks, then accept the message. This allows relaying based on authentication.

Check the client's IP address against an explicit blacklist table. Do not put any 'OK' entries in this table or else you will become an open relay for the given machine.

Permit machines on the local network to send, if the above sanity checks pass. This uses the IP address as a form of authorization.

Reject any HELO/EHLO parameter that contains my own domain name or IP address. The HELO means "Hello, my name is <blah>". No one else has my name (or IP address), and therefore legitimate mail won't be using my own name.

Reject any invalid or not fully-qualified host name in the HELO parameter. This particular check can have some collateral damage, if the sender's software does not follow the rules of SMTP. It stops a lot of spam, though, and any quality mail server software run by a competent (or teachable :-)) admin

won't be blocked by this test.

For senders in one of the "freemail" domains (hotmail, msn, aol, yahoo), verify that the host sending the mail is also in that domain. Those services always send mail out using their own systems. This catches a lot of forged senders on spam.

Reject "unauthorized" destinations. A destination is authorized if it is for a domain that I handle locally or one I relay for as a backup server. All others are attempts by spammers to use my machine to pass on their junk to others. (note that authorized relaying is allowed much earlier with the authenticated and mynetworks checks)

Reject if recipient doesn't exist. This is very important to manage the load on your mail server. If you instead accept all local recipients and then bounce the ones that don't exist (like qmail does, and some others can do if so configured) then you will end up with a lot of undeliverable bounces stuck in the queue slowing down the system.

2.3 Content matching

After all of the SMTP-level tests have passed, I perform some simple pattern matches on the data of the message. The lines in postfix' main.cf are

```
header_checks = pcre:/etc/postfix/header_checks
mime_header_checks = $header_checks,
                     pcre:/etc/postfix/mime_header_checks
body_checks = pcre:/etc/postfix/body_checks
strict_8bitmime = yes
```

2.3.1 header_checks explained

As I look through the actual contents (shown below) of my header_checks file I see two sets of checks. One set is to reject, during the SMTP session, the "you have a virus" messages. I reject these because I most certainly do not have a Windows virus or worm on my Debian system. I get such alerts due to poor logic in the virus scanners various sites have deployed – they

either arrive via a mailing list or because a worm spoofed my address in it somewhere (envelope or headers). The second set is to remove headers that I feel are pointless. They serve no purpose to me, and I figure as long as I can remove them so easily, why waste the disk space storing them? Note: the reasoning I use may or may not apply to your user base. That is for you to decide. It works for me (my userbase is me) :-).

2.3.2 body_checks explained

I have pflogsumm[4] run from a cron job to email me a summary of the mail server's load. Naturally, since it includes a summary of rejects, it would be blocked by the patterns below. The first test is to see if the text looks like a pflogsumm report and if so to skip the rest of the tests on that line.

Below that I discard anything that looks like a Windows executable (or RIT's not-all-that-helpful "I replaced the virus with this alert" message). I discard, rather than reject, because almost always the sender address is spoofed and it is not good to Joe-job[6] an innocent bystander. Note that I don't try to parse MIME-encoded filenames and create a comprehensive list of "potentially dangerous" file names. Testing for the (base64-encoded) executable content is simpler and more effective. Yes, this means that no one can email me a windows executable directly. However, I would have little use for such content on my Debian system and furthermore, if you actually had a legitimate reason to send me one you could easily package it in a tarball or zip file.

Finally, I've noticed a sufficiently annoying number of messages that self-proclaim to be virus free. How utterly pointless! Do you think a virus will announce itself? Of course not. Anyone can put a "this message was scanned" tag in a message regardless of whether or not is was scanned. Furthermore, my system is not vulnerable to such malware. This is why I remove such noise from the data stream in the last few patterns.

2.4 Sampling of my Table's Contents

postfix is very table-driven. This is a rather neat design, and makes certain aspects of configuration very simple and straightforward. Anyways, some of the contents of various tables are rather important to the effectiveness of my anti-spam measures. They, and a little bit extra, are revealed here.

2.4.1 access-helo.pcre

```
/^[0-9.]+$/                 550 Your software is not RFC 2821 compliant
```

2.4.2 access-helo

```
localhost.localdomain   REJECT  localhost? Really?  Nah, fix your 'hosts' file!
dman13.dyndns.org       REJECT  Don't use my own hostname
66.66.61.253            REJECT  Don't use my own IP address
dman.ddts.net           REJECT  Don't use my own hostname
```

2.4.3 access-freemail-client

Note that this implementation is really an approximation of the intended check. For example, this test allows an @aol.com sender to come from a machine on earthlink.net's network. However, in practice it is close enough to the intended test and many times simpler than a completely proper implementation would be. (a proper implementation would have a separate table for each of the domains in question and would take into account the hotmail–msn relationship)

```
yahoo.com           OK
hotmail.com         OK
msn.com             OK
aol.com             OK
earthlink.net       OK
excite.com          OK
excitenetwork.com   OK
```

2.4.4 header_checks

```
/^Subject:[ ]*Infected E-Mail$/       REJECT Bogus virus warning detected [110].
                                             Contact <postmaster> for details.
/^From: NAV for Microsoft Exchange/   REJECT Bogus virus warning detected [102].
                                             Contact <postmaster> for details.
/^From: F-Secure Anti-Virus for Internet Mail/
                                      REJECT Bogus virus warning detected [111].
                                             Contact <postmaster> for details.
/^Subject: .*(?:NAV|Norton AntiVirus) detected (?:and quarantined )?a virus/
                                      REJECT Bogus virus warning detected [103].
                                             Contact <postmaster> for details.
/^Subject: .*ScanMail for Lotus Notes/ REJECT Bogus virus warning detected [104]
                                              Contact <postmaster> for details.
/^Subject: .*Symantec AVF detected a.*virus/
                                      REJECT Bogus virus warning detected [105].
                                             Contact <postmaster> for details.
/^Subject: .*Virus Alert/             REJECT Bogus virus warning detected [106].
                                             Contact <postmaster> for details.
/^Subject: .*A Virus was detected/    REJECT Bogus virus warning detected [107].
                                             Contact <postmaster> for details.
/^Subject: .*VIRUS IN YOUR MAIL/      REJECT Bogus virus warning detected [108].
                                             Contact <postmaster> for details.
/^Subject: .*Virus Detected by Network Associates/
                                      REJECT Bogus virus warning detected [109].
                                             Contact <postmaster> for details.
|^X-Mailer: ravmd/8.3.2|  REJECT  Mail from virus scanners is not accepted [100]
/^X-Mailer: MailScanner/  REJECT  Mail from virus scanners is not accepted [112]
/^X-[^:]*MailScanner: Found to be infected/
              REJECT  Joe-Jobbing is an unacceptable abuse of this system. [101]
/^X-Auto-Generated: McAfee antivirus plugin/
                          REJECT  Mail from virus scanners is not accepted [113]

# Pointless.
/^X-Virus-Scanned:/     IGNORE
/^X-AntiVirus:/         IGNORE
/^X-RAVMilter-Version:/ IGNORE
```

```
/^X-Mailscanner[^:]*:/  IGNORE
/^X-[^:]*MailScanner:/  IGNORE
/^X-Kaspersky-Antivirus:/  IGNORE
/^Thread-Topic:/        IGNORE
/^X-Sun-Charset:/       IGNORE

/^X-MIMEOLE:/           IGNORE
/^X-MSMail-Priority:/   IGNORE
```

2.4.5 body_checks

```
# Hmm, the pflogsumm report will end up including text rejected below
/^ {6,11}\d{1,6}[ km] /    OK

# RIT has already identified the junk
/^The RIT mail system has detected an attachment that is considered high risk$/
                                                                    DISCARD

# All .exe files from MSVC have the same starting bytes
/^TVqQAAMAAAAEAAAA\/\/8AALgAAAAAAAAQAAAAAAAAAAAAAAAAAAAAA[A:]A*$/
                                                  DISCARD  MSVC executable

# (anchor sufficiently to avoid rejects passing on the patterns)
/^-----------+ +(?:Sify )?Virus Warning Message /
                                       REJECT  Bogus virus warning detected [300].
                                               Contact <postmaster> for details.
# Joy.  Multiple languages.
/^------ Message du Moteur Antivirus /
                                       REJECT  Bogus virus warning detected [305].
                                               Contact <postmaster> for details.
/Antigen for Exchange found[^\/[]/   REJECT  Bogus virus warning detected [301].
                                               Contact <postmaster> for details.
/^Sophos Plc MailMonitor for Domino/ REJECT  Bogus virus warning detected [302].
                                               Contact <postmaster> for details.
/^--- Dr\.Web report ---/            REJECT  Bogus virus warning detected [303].
                                               Contact <postmaster> for details.
/^The scanned e-mail has your address in the <From> header field./
```

```
                                        REJECT  Bogus virus warning detected [304].
                                                Contact <postmaster> for details.
/^------------ Trend GateLock/          REJECT  Bogus virus warning detected [306].
                                                Contact <postmaster> for details.
/^WAKWAK Virus Detect System has found the file which is/
                                        REJECT  Bogus virus warning detected [307].
                                                Contact <postmaster> for details.


# Dumb.  Just plain dumb.
/^Outgoing mail is certified Virus Free\.$/                                 IGNORE
|^Checked by AVG anti-virus system \(http://www\.grisoft\.com\)\.|         IGNORE
#Version: 6.0.510 / Virus Database: 307 - Release Date: 8/14/03
#Version: 6.0.520 / Virus Database: 318 - Release Date: 18/09/2003
#   Note: some variants have a different date format
|^Version: .\..\.... / Virus Database: ... - Release Date: .?./../..(?:..)?$|
```

3 maildrop, spambayes

maildrop is handling local delivery for me. It's filter file pipes each message through spambayes to be classified by a variation of the Bayesian method. Any messages that spambayes marks as spam are automatically discarded. Any marked as "unsure" are filed in the "unsure" folder for manual checking (and continued training of the classifier) and any marked as non-spam are delivered as usual to the appropriate folder.

The key to effective Bayesian classification is an accurate statistical history. This history is created by manually (and correctly!) selecting a set of span and non-spam messages and telling the classifier which is which. While this need for training often deters people from using it, and is what delayed my experimentation with it, I discovered that even with very little training spambayes was quite accurate and effective. The training never ceases, but since very few messages are not correctly tagged I don't train very much any more. FWIW, my current (at the time of this writing) database consists of stats from 377 ham and 376 spam messages with a total of 33839 distinct "words".

References

[1] postfix: http://www.postfix.org/

[2] maildrop: http://www.flounder.net/∼mrsam/maildrop/

[3] spambayes: http://spambayes.sourceforge.net/

[4] pflogsumm: http://jimsun.linxnet.com/postfix_contrib.html

[5] Jim Seymour's postfix anti-UCE "cheat sheet" http://jimsun.linxnet.com/misc/postfix-anti-UCE.txt

[6] Joe-job (definition): http://members.cox.net/joejob/#WhatIsAJoeJob