
SpamAssassin

SpamAssassin is an extensible email filter which is used to identify spam.

Yung-Zen Lai (yzlai@tp.edu.tw)

2004/10

Agenda

- n Introducing SpamAssassin
 - q How It Works
 - q Mailer and SpamAssassin
 - n SpamAssassin Basics
 - q Prerequisites
 - q Building SpamAssassin
 - q Invoking SpamAssassin
 - n SpamAssassin Rules
 - n SpamAssassin as a Learning System
 - q AutoWhiteListing
 - q Bayesian Filtering
-

Agenda (cont.)

- n Integrating SpamAssassin with Postfix
 - q Postfix Architecture
 - q Spam-Checking During Local Delivery
 - q Spam-Checking All Incoming Mail
 - n AMaViSd.conf
 - n Using SpamAssassin as a Proxy
 - q Using Pop3proxy
 - q Using SAproxy Pro
-

Introducing SpamAssassin

- n Software for analyzing email messages, determining how likely they are to be spam, and reporting its conclusions.
 - q Uses a large number of different kinds of rules and weights them.
 - q Easy to tune the scores associated with each rule or to add new rules.
 - q Can adapt to each system's email environment, learning to recognize which senders are to be trusted and to identify new kinds of spam.
-

From: "xbwhc" <acfxi@ms73.hinet.net>

Reply-To: "xbwhc" <acfxi@ms73.hinet.net>

To: netadm@tp.edu.tw

X-Mailer: punk digestible egregious

Subject: =?ISO-8859-1?Q? *SPAM-Mail*

XYZ=B8=EA=B0T=A4u=A7{=A7=F3=B7s=A7=D6=B3=F8chosen?=
^

Date: Mon, 18 Oct 2004 22:02:19 +0400

MIME-Version: 1.0

Organization: punk digestible egregious

Content-Type: multipart/alternative; boundary="====85167443147683=_"

X-Amavis-Alert: BAD HEADER Non-encoded 8-bit data (char B8 hex) in message header 'Subject' Subject:

XYZ\270\352\260T\244u\247{\247\363\267s\247\326\263\370... ^

X-Spam-Status: Yes, hits=15.9 tagged_above=5.0 required=8.0
tests=BAYES_99, HTML_FONTCOLOR_RED, HTML_MESSAGE,
RAZOR2_CF_RANGE_11_50, RAZOR2_CHECK, RCVD_IN_DSBL,
RCVD_IN_NJABL, RCVD_IN_NJABL_DIALUP, RCVD_IN_RFCI,
SUBJ_ILLEGAL_CHARS

X-Spam-Level: *****

X-Spam-Flag: YES

How SpamAssassin Works

- n Check consistency and adherence to Internet standards. (RFC 2821, 2822, ...)
 - n Check headers and body for phrases or messages elements commonly found in spam.
 - n looked up headers and body in several online databases that track message checksums of verified spam messages.
 - n Lookup up the sending system's IP address in several online lists of sites that have been used by spammers or are otherwise suspicious. (RBL)
-

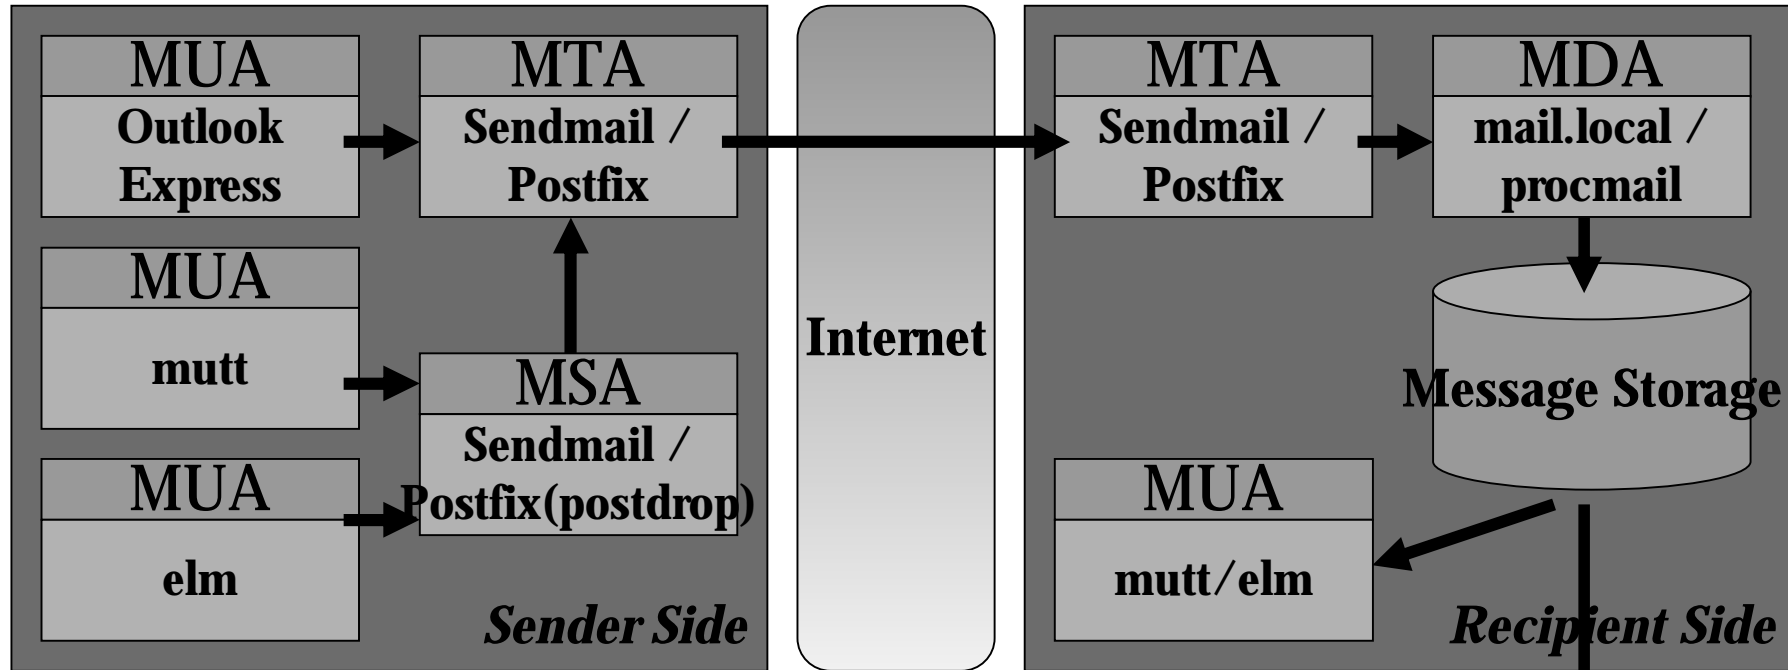
How SpamAssassin Works (cont.)

- n Specific addresses, hosts, or domains can be blacklisted or whitelisted.
 - n Can be trained to recognize the types of spam that you received.
 - n The sending system's IP address can be compared to the sender's domain name using the Sender Policy Framework (SPF) protocol. (version 3.0)
 - n Can privilege senders who are willing to expend some extra computational power in the form of Hashcash. (version 3.0)
-

Mailers and SpamAssassin

- n Typical mail transmission
 - n Scanning at the MTA
 - n Scanning at the MDA
 - n Scanning with a POP Proxy
-

Message flow




UA: User Agent

TA: Transfer/Transport Agent

DA: Deliver Agent

SA: Submission Agent

AA: Access Agent



SpamAssassin Basics -- Prerequisites

- n Written in Unix or Unix-Like environment that includes Perl 5, preferably 5.6.1 or later.
 - n Up-to-Date versions of the following Perl modules
 - q ExtUtils::MakeMaker
 - q File::Spec
 - q Pod::Usage
 - q HTML::Parser
 - q Sys::Syslog
 - q DB_File
 - q Digest::SHA1
 - q Net::DNS
-

Building SpamAssassin

- n perl -MCPAN -e shell
 - q o conf prerequisites_policy ask
 - q install Mail::SpamAssassin
 - n Build from source tarball
 - q <http://spamassassin.apache.org>
 - q perl Makefile.PL
 - n Linux: Download the source RPM
 - q <http://spamassassin.apache.org>
 - n FreeBSD: Build from the ports collection
 - q /usr/ports/mail/p5-Mail-SpamAssassin-3.xx
-

Basic Configuration & Test

- n [/usr/local] /etc/mail/spamassassin/local.cf
 - q required_hits/required_score
 - q report_safe
 - q rewrite_subject (2.x only)
 - q rewrite_header (Subject, From, or To headers 3.0 only)
 - q skip_rbl_checks
 - n Test
 - q spamassassin --test-mode < sample-spam.txt
 - q spamassassin --test-mode < sample-nospam.txt
 - n Untagging
 - q spamassassin --remove-markup < marked-message >
unmarked-message
-

Invoking SpamAssassin

n With procmail

q Scanning at the MDA (/etc/procmailrc)

```
DROPPRIVS=yes
PATH=/bin:/usr/bin:/usr/local/bin
SHELL=/bin/sh

# Spamassassin
:0fw
< 300000
|/usr/bin/spamassassin
```

Invoking SpamAssassin (cont.)

- n Using spamc/spamd
 - q Client/Server Architecture
 - q spamc is a lightweight client, written in C

```
$ /usr/bin/spamd --daemonize --pidfile /var/run/spamd.pid
```

```
$ spamc -c < sample-nospam.txt
```

```
0.0/5.0
```

```
$ spamc -c < sample-spam.txt
```

```
1000.0/5.0
```

Invoking SpamAssassin (cont.)

- n Invoking spamc with procmail (procmailrc)
 - q Just as spamc is run in place of the spamassassin

```
DROPPRIVS=yes
PATH=/bin:/usr/bin:/usr/local/bin
SHELL=/bin/sh

# Spamassassin
:0fw
< 300000
|/usr/bin/spamc
```

Invoking SpamAssassin (cont.)

- n spamd caches the spam-checking rules in memory when it starts up.
- n The daemon must be signaled whenever changes the SpamAssassin rulesets or systemwide configuration file.
- n spamd reloads configuration files when it receives a HUP signal.

```
$ kill -HUP `cat /var/run/spamd.pid`
```

Invoking SpamAssassin (cont.)

n In a Perl Script

```
#!/usr/bin/perl

use Mail::SpamAssassin;

my @lines = <STDIN>;
my $mail = Mail::SpamAssassin::NoMailAudit->new(data => \@lines);
my $spamtest = Mail::SpamAssassin->new();
my $status = $spamtest->check($mail);
$status->rewrite_mail() if ( $status->is_spam() );
print $status->get_full_message_as_text();
```

SpamAssassin and the End User

- n True Negatives(ham)

- q Both you and SpamAssassin agree are non-spam

- n True Positives(spam)

- q Both you and SpamAssassin agree are spam.

- n False Positives

- q SpamAssassin marks a message as spam that you actually wanted to receive.

- n False Negatives

- q A missed spam.

SpamAssassin Rules

- n All static tests
 - q Don't change over time
 - n Pattern-based & Network-based tests
 - n How they are written and scored
 - n How to modify the score of a build-in test
 - n How to write our own custom tests
 - n Whitelist and Blacklist rules
-

A test definition and score

header	FROM_STARTS_WITH_NUMS	From =~ /^d\d/
describe	FROM_STARTS_WITH_NUMS	From: starts with nums
score	FROM_STARTS_WITH_NUMS	0.390 1.574 1.044 0.579

- n The header directive defines it as a test that will be applied to the message headers and give the test name (FROM_STARTS_WITH_NUMS)
 - n The describe directive provides a human-readable description
 - n SpamAssassin will add 0.390, 1.574, 1.044, or 0.579 to the spam score with network and Bayesian tests are in use or not in use
-

Modifying the Score of a test

- n Adjustment scores in systemwide configuration file
 - q `/etc/mail/spamassassin/local.cf`
 - q `score HTML_WIN_OPEN 2`

 - n Write our own systemwide tests
 - q Regular expression matching
 - q Eval test defined in `Mail::SpamAssassin::Evaltests`
-

Header Tests

- n Existence of a header
 - q header TESTNAME exists:headername
 - n Test single header
 - q header TESTNAME headername =~ /regexp/modifiers
 - n Test single header not match
 - q header TESTNAME headername !~ /regexp/modifiers
 - n Test multiple header
 - q header TESTNAME =~
/^(:headername1|headername2|...):.*modifiers
 - n All capital letter
 - q header SUBJ_ALL_CAPS eval:subject_is_all_caps()
 - n RBL
 - q header TESTNAME =~ eval:check_rbl('ZoneID', 'Server')
 - q check_rbl_txt()
 - q check_rbl_sub()
-

Body Tests

- n The body, rawbody, full directives
 - q Can be tested against a regular expression
 - q Can be submitted to an eval test
 - n The body directive appears to a person reading the message in an text-based MUA
 - n The rawbody directive appears to a person reading the message in an HTML-based MUA
 - n The full directive defines a test to be applied to the full text of a message
-

URI Tests

- n Test on all URIs that appear in an email message
 - q uri MAILTO_TO_REMOVE /^mailto:.*?remove/is
- n SpamAssassin 3.0 includes a plug-in to test the uridnsbl directive

uridnsbl	URIBL_SBLXBL	sbl-xbl.spamhaus.org. TXT
header	URIBL_SBLXBL	eval:check_uridnsbl('URIBL_SBLXBL')
describe	URIBL_SBLXBL	Contains a URL listed in the SBL/XBL
score	URIBL_SBLXBL	10

Test Examples

- n For a To, From, or Cc header
 - q header XXX ALL =~ /^(?:to|cc|from):.*abc@def/im
 - n For the existence of the X-PMFLAGS header
 - q header XXX exists:X-PMFLAGS
 - n For long lines of hexadecimal code in the message body
 - q header XXX /[0-9a-fA-F]{70,}/
 - n For Subject header in all capital letters
 - q header XXX eval:subject_is_all_caps()
 - n For a message that include HTML to open a new window
 - q body XXX eval:html_test('window_open')
 - n For an HTTP URI in the message
 - q uri XXX /^http:ooxxabcd/is
-

The Built-in Tests

- n Over 700 test rules

- q /usr/share/spamassassin

- q <http://spamassassin.apache.org/tests.html>

- n 10_misc.cf

- q Defines special rules that are not spam tests

- q Include templates for the spam report

- n 20_fake_helo_tests.cf

- q Defines a set of rules that use the eval test
check_for_rdns_helo_mismatch()

Whitelists and Blacklists

n Whitelisting senders

q whitelist_from yzlai@hotmail.com

q whitelist_from *@tp.edu.tw

n Whitelisting senders by relay

q whitelist_from_rcvd *@tp.edu.tw smtp.tp.edu.tw

n Whitelisting recipients

q whitelist_to netadm@tp.edu.tw

n Blacklist

q blacklist_from netadm@tp.edu.tw

q blacklist_to spamtrap@tp.edu.tw

SpamAssassin as a Learning System

- n Learns each sender's history of sending spam or non-spam messages
 - n Modifies the spam score of their subsequent mailings on the basis of this history
 - n Trying to reduce false positives
 - n Can also reduce false negatives, but this happens infrequently enough.
-

Autowhitelisting

- n SpamAssassin maintains a database keyed on message senders' email address and the IP addresses of their nearest untrusted relay.
 - n The average sender score – the total score divided by the number of messages received – is used to modify the spam score of new messages from that sender
 - n Use *auto_whitelist_factor* directive to set the multiplier between a message's spam score and the sender's history average score.
-

Autowhitelisting (cont.)

Message number	Message score (before)	Sender average score	Score after autowhitelist with give AWF			
			0	.5	.75	1
1	2	(none)	2	2	2	2
2	1	2	1	1.5	1.75	2
3	1	1.5	1	1.25	1.375	1.5
4	0	1.33	0	0.67	1.00	1.33
5	2	1.0	2	1.5	1.25	1.0
6	6	1.2	6	3.6	2.4	1.2

Using Autowhitelists in Perl

```
#!/usr/bin/perl

use Mail::Spamassassin;

my @lines = <STDIN>;
my $mail = Mail::SpamAssassin::NoMailAudit->new(data => \@lines);
my $spamtest = Mail::SpamAssassin->new();
my $awl = Mail::SpamAssassin::DBBasedAddrList->new();
$spamtest->set_persistent_address_list_factory($awl);
my $status = $spamtest->check($mail);
$status->rewrite_mail() if ( $status->is_spam() );
print $status->get_full_message_as_text();
```

Bayesian Filtering

- n Bayesian filtering is based on Bayes' Theorem, a statement of probability theory
 - n Can reduce both false positives and false negatives
 - n The probability of some event given a test result depends on the baseline probability of the event before the test result is known and on the discriminating power of the test
 - n The more the test result can increase (or decrease) the probability from baseline, the stronger the test.
 - n Some are patterns that only occur in spam (and never in non-spam) ...
-

Bayesian Filtering (cont.)

n Configuration

- q use_bayes
 - q bayes_auto_learn
 - q bayes_auto_learn_threshold_nonspam
 - q bayes_auto_learn_threshold_spam
 - q bayes_ignore_from address
 - q bayes_ignore_to address
 - q bayes_auto_expire
 - q bayes_expiry_max_db_size
-

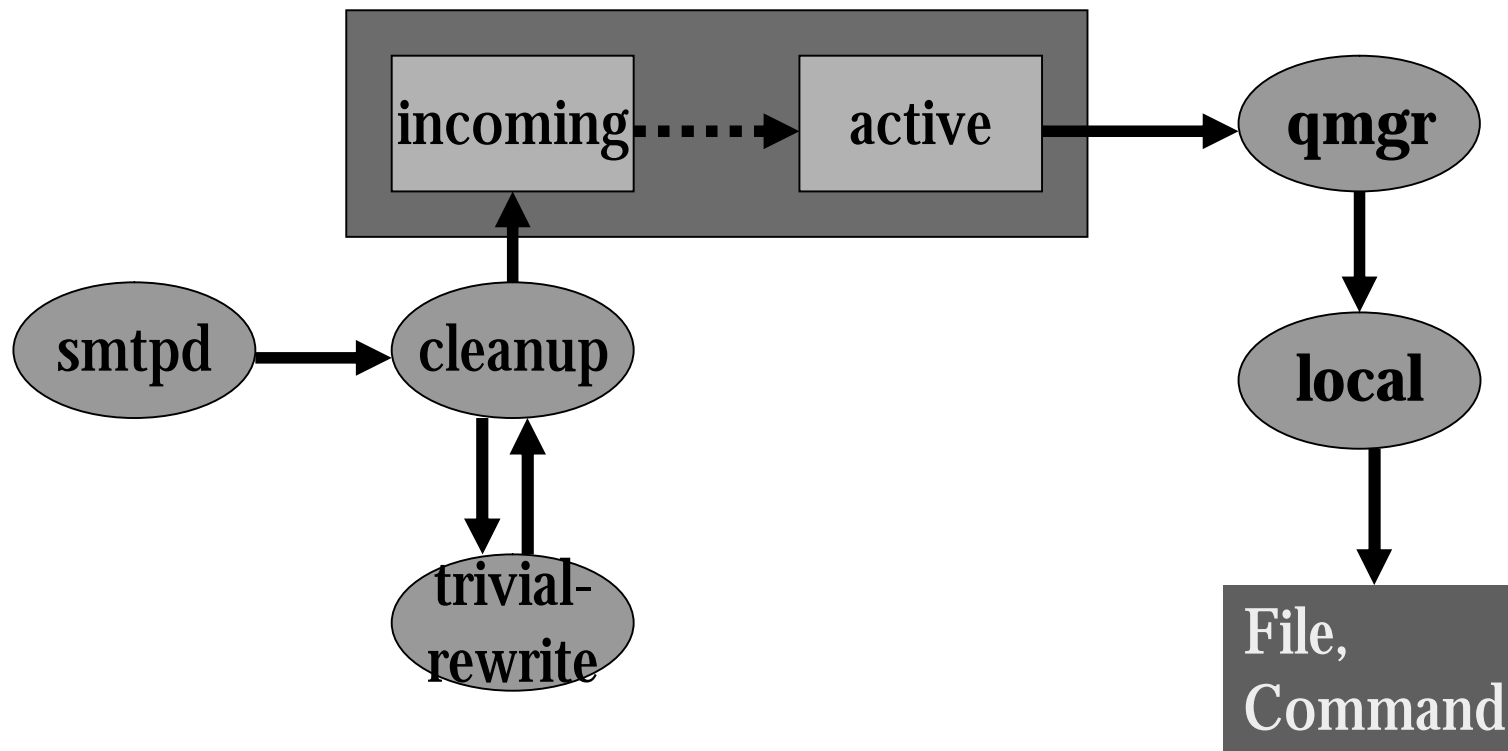
Bayesian Filtering (cont.)

n Training

- q `sa-learn --mbox --spam Mail/spam`
 - q `sa-learn --mbox --ham mail/nospam`
 - q `sa-learn --mbox --spam /var/spool/mail/spamtrap`
 - q `sa-learn --dir --spam ...`
 - q `sa-learn --no-rebuild --spam ...`
 - q `sa-learn --rebuild`
-

Integrating with Postfix

n Postfix architecture



Integrating with Postfix (cont.)

n Checking During Local Delivery

- q Easy to set up
- q Run *spamd*, and the *procmail* recipe can use *spamc* for faster spam-checking
- q User preference files, autowhitelisting, and Bayesian databases can be used

n Postfix's mail.cf

- q `mailbox_command = procmail`
-

Integrating with Postfix (cont.)

- n Checking All Incoming Mail
 - q Postfix's *content_filter* parament
 - q To a Program
 - q To a Daemon
 - n More effective
 - n Building a Spam-Checking Gateway
-

Using a Program as a Content Filter

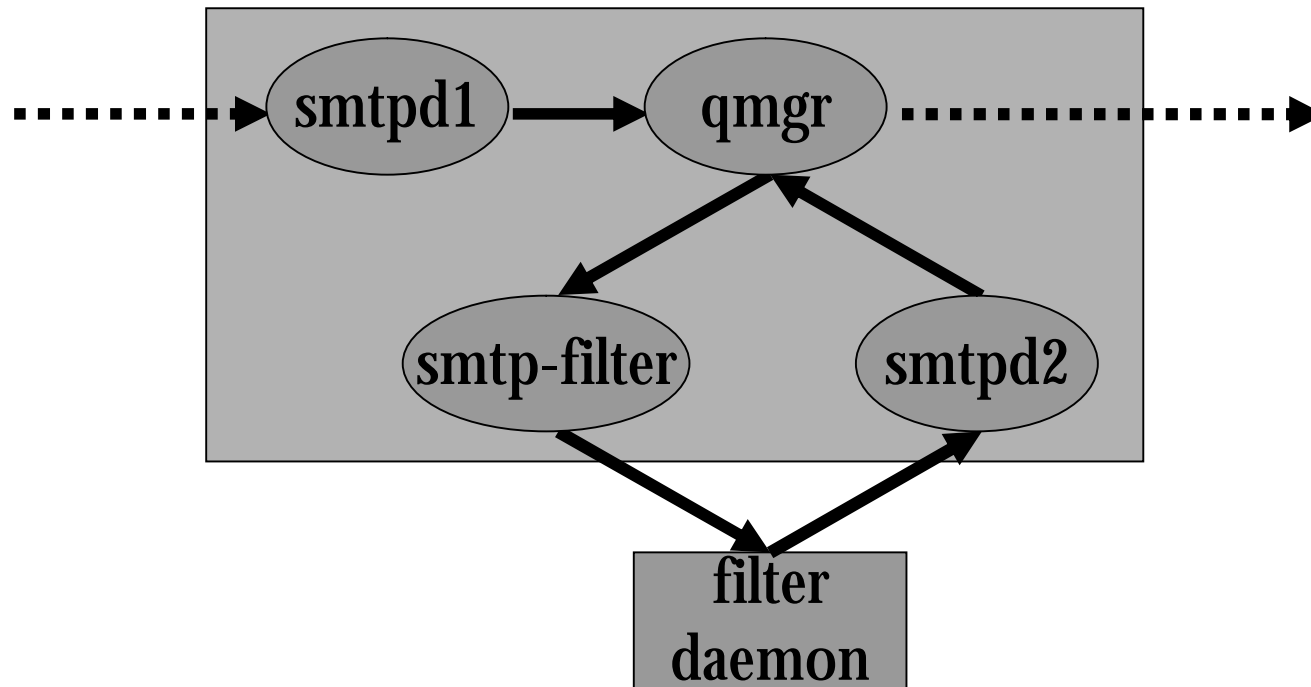
```
#!/bin/sh

sender=$1
shift
recip="$@"
if [ "$#" -eq 1 ]; then
    /usr/bin/spamc -u $recip
else
    /usr/bin/spamc
fi | /usr/sbin/sendmail -i -f $sender -- $recip
exit $?
```

```
n main.cf
  q content_filter = spamcheck:
n master.cf
  q smtpcheck unix - n n - - pipe
    flags=Rq user=spamfilt argv=/usr/local/bin/pf-spamfilt ${sender} ${recipient}
```

Using a Daemon as a Content Filter

- q Daemon-based filtering
 - n more efficient



Using a Daemon as a Content Filter (cont.)

n Install a daemon that performs content-filtering

q amavisd-new

n main.cf

```
content-filter = smtp-filter:[127.0.0.1]:10024
```

n master.cf

```
smtp-filter  unix  -  -  n  -  10  smtp
```

```
-o myhostname=localhost
```

```
localhost:10025  inet  n  -  n  -  10  smtpd
```

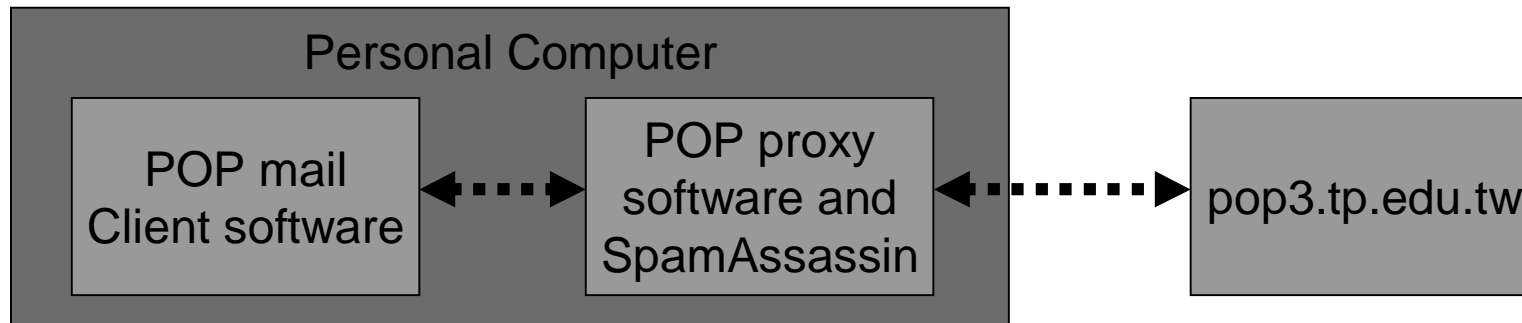
```
-o content_filter=
```

AMaViSd.conf

```
n $mydomain = 'tp.edu.tw';  
n $forward_method = 'smtp:127.0.0.1:10025';  
n $max_servers = 15;  
n $final_spam_destiny = D_DISCARD;  
n $sa_tag_level_deflt = 5.0;  
n $sa_tag2_level_deflt = 8.0;  
n $sa_kill_level_deflt = 20.0;  
n $sa_spam_subject_tag = '*SPAM-Mail*';  
n $sa_... = ...;
```

Using as a Proxy

n Message flow



Using as a Proxy (cont.)

n Using Pop3proxy

- q Download and Extract Pop3proxy from <http://mcd.perlmonk.org/pop3proxy/>
 - q Install Perl for Windows which includes Time::HiRes module
 - q Download and Extract SpamAssassin from <http://spamassassin.apache.org/>
 - q ``$ perl pop3proxy.pl --host pop3.tp.edu.tw``
 - q Reconfigure the POP Client
 - n Connect to *localhost* instead of the usual POP server
-

Using as a Proxy

n Using SAproxy Pro

- q A commercialized version of Pop3proxy

- q <http://www.statalabs.com/>

- q Reconfigure the POP Client

- n Connect to *localhost* instead of the usual POP server

- n <account>:<pop server> in Account Name

The End

- n Reference

- q *SpamAssassin – The Open Source Solution to SPAM*, Alan Schwartz, O'Reilly, July 2004

- n Thank you!
