Global Information Assurance Certification Paper

Anti Spamming – How to filter unsolicited e-mail on your mail server

Nam Tran
December 27, 2001

This document is written from real work experience. In this document, the main software components, sendmail™ and IP Filter will be discussed. The complete solution would require a combination of additional software and hardware components. The solution herein is for the Internet Service Providers (ISPs) and organizations providing email service to their employees, rather than for the end users.

Understanding Spam

What is spam?

> *"Spamming is the scourge of electronic-mail and newsgroups on the Internet. It can seriously interfere with the operation of public services, to say nothing of the effect it may have on any individual's email mail system. ... Spammers are, in effect, taking resources away from users and service suppliers with out compensation and without authorization."*
>
> Vint Cerf, Senior Vice President, MCI Worldcom
> And acknowledged "Father of the Internet"
> (quoted from http://www.cauce.org/about/problem.shtml)

Spam is the junk mail of the Internet, sending with an attempt to force the message on people who would not otherwise choose to receive it. Typically, spam is mass mailing of advertisements and promotions to large mailing lists, also called Unsolicited Commercial Email, or UCE. While UCE is often sent to large mailing list, often an application is used to walk the alphabet on a known email gateway. Spam costs the sender very little to send; but it could cost the receiving IPSs a lot of money. Servers receiving the mailings are overwhelmed by the sheer volume, which may lead to a "denial of service". This results in a loss of legitimate email service to your employees or customers.

According to CAUCE (Coalition Against Unsolicited Commercial Email), its research on SPAM-L mailing lists (URL: http://www.claws-and-paws.com/spam-l/) and USENET newsgroup in the news.admin.net-abuse.* hierarchy, the most commonly seen UCEs are:

- Chain letters
- Pyramid schemes (including Multilevel Marketing, or MLM)
- Other "Get Rich Quick" or "Make Money Fast" (MMF) schemes
- Offers of phone sex lines and ads for pornographic web sites
- Offers of software for collecting email addresses and sending UCE
- Offers of bulk emailing services for sending UCE

- Stock offerings for unknown start-up corporations
- Quack health products and remedies
- Illegally pirated software ("Warez")

Spam is difficult to fight because of the sheer volume. AOL has estimated that up to 30% of incoming email on its networks is spam. Spam begins when mailing lists are sold to companies who want to instant mailing opportunities. With one push of a button, spammers can flood networks with advertisements. According CAUCE, with a PC and a dialup connection, a spammer can send hundreds of thousands of messages per hour. Although the spammer pays for mailing list, ISPs pay for the network use, and that cost eventually trickles down to the customer. ISPs also pay in that customer satisfaction decreases and customer complaints increase.

Why is spam bad?
Usually users get upset when they receive unsolicited emails. Spam is bad for the recipients because of several reasons:

1. Greater cost (money and time) for the recipients than for the sender
2. Fraud – trick the recipients to open the junk mail, for example, by making the mail "subject" look like it is anything other than an advertisement.
3. Theft of other's resources – spammers sending their mail via innocent intermediate systems.
4. It's all garbage – it's all "get rich quick", "make money fast" and the like, etc.
5. It's annoying – wait so long to retrieve mails, and when it's done, all you get is what you don't want to see.
6. It might be illegal – some forms of spam (such as child pornography) are illegal in some countries on the Internet.

You can find greater details about this topic on http://spam.abuse.net/spam/spambad.html and http://www.cauce.org/about/problem.shtml.

How To Fight Spam

Filtering is one method most ISPs use to fight unwanted email. Lists of commercial e-mailers exist, and algorithms can determine if email matches the bulk email profile, and this mail can then be blocked. Unfortunately, spammers have ways to work around blocks and filters. According to CAUCE:

> "One of the most common tricks is to relay their messages off the mail server of an innocent third party. This tactic doubles the damages: both the receiving system, and the innocent relay system are flooded with junk email. And for any mail that gets through, often times the flood of complaints goes back to the innocent site because they were made to look like the origin of the spam. Another common trick that spammers use is to forge the headers of messages, making it appear as though the message originated elsewhere, again providing a convenient target." (http://www.cauce.org/about/problem.shtml, page 3)

Nonetheless, filters and blocks are still the most viable solution, even though this solution will cost an IPS money, and takes up CPU and processing time. One of the most effective ways to filter spam is implementing a sendmail™ Mail Transport Agent (MTA), coupled with MAPS Realtime Blackhole Lists, MAPS Relay Spam Stopper and the MAPS Dial-up User List, running on a Sun Solaris platform, preferably Solaris 8.

Solution Overview

Spam filtering can be done via:

- IPFilter
- Request server not in DNS
- Relay-Domains
- Access.db
- MAPS[sm]

On the above list, IPFilter is a free software package. The other four bullet items can be implemented using sendmail™.

The solution:

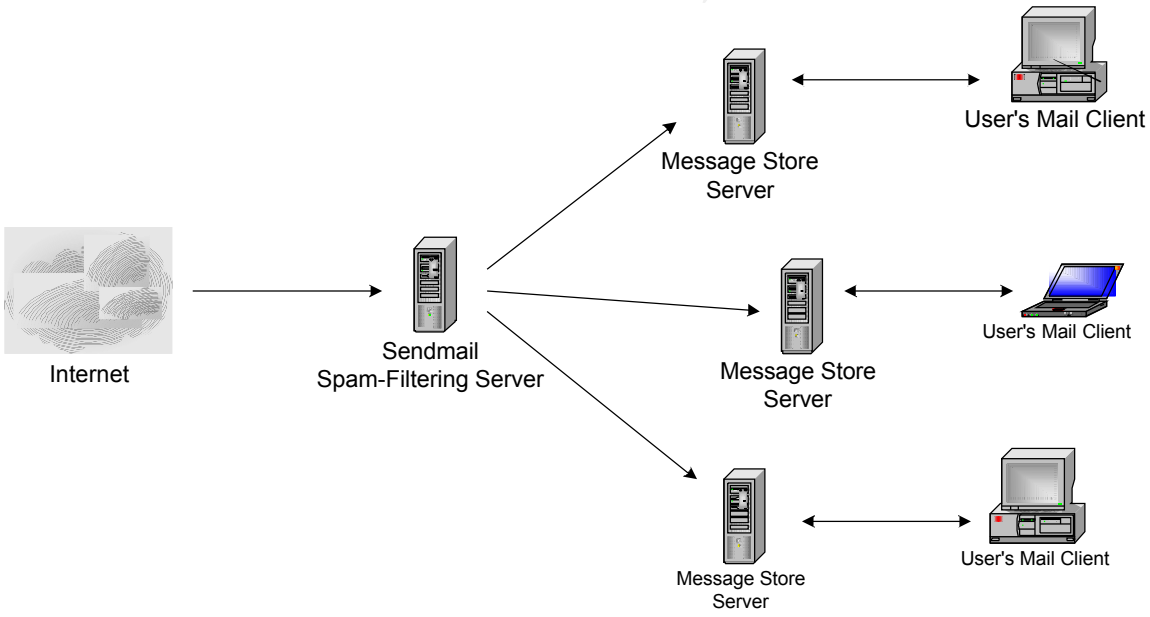

Inbound Mail Filtering

As presented on the picture above, a sendmail™ server is put in front of all customer-facing mail servers. The sendmail server is of course loaded with spam filtering schemes. The goal of the sendmail™ server is to prevent unnecessary load on the customer-facing mail/messaging servers (SMTP and POP), and to improve performance for the customer. sendmail™ gives operations the ability to have tools to combat unnecessary load on that server. The sendmail™ server accomplishes these goals by:

- Changing the DNS so that incoming email is redirected to the sendmail™ server, and not to the customer-facing servers, as portrayed in the picture above.
- Adding new MX records for each customer-facing mail server that will be serviced bye this spam-filtering server.
- Configure the router(s) attached to the messaging server to prevent SMTP connections to the customer-facing mail server on TCP port 25 unless the connection comes from local IP space.

Although this document does not specifically describe how the router should be configured, a pseudo-description of the commands used might be:

Allow INBOUND TCP Port 25	(from Network to User)
Deny ALL Other TCP Port Connections	(from Network to User)
Allow ALL Outbound TCP Port Connections	(from User to Network)

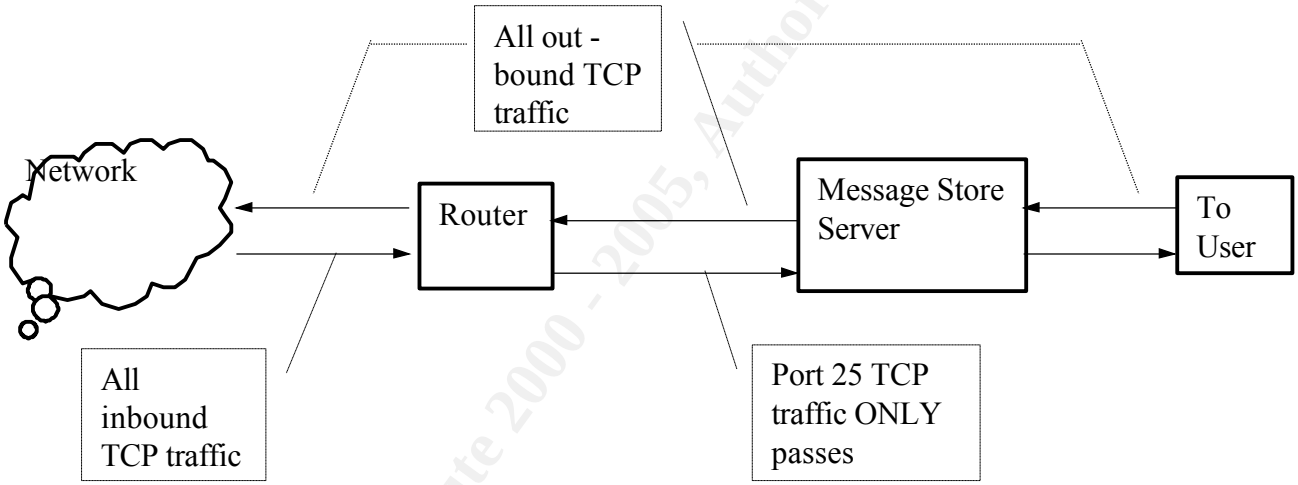

Routers filter all Simple Mail Transfer Protocol (SMTP) traffic and protect it from customers and the outside world. The only SMTP traffic allowed to the customer-facing messaging servers is from hosts with an MX record for xyz.com, or from the new system with the MX record for <subdomain>.xyz.com.

Why Sendmail™?

sendmail™ version 8.11.4 and later contains a number of Spam-prevention features. These include:

- The use of he *helo/ehlo* commands
- The ability to refuse *vrfy* and *expn* requests
- DNS verification of the hostname given in the 'MAIL from:' command
- Comparison of the IP address or domain name of a relaying host with a database of blocked users and blocked sites
- Comparison of the 'MAIL from:' address with a database of blocked users and sites
- Blocking relaying of mail to other ISPs.

- Rejecting messages based on IP address, domain, email address, to/from header and subject line

Also, sendmail™ is able to perform the following operations:

- Redirect mail to the appropriate subdomain, if any, such as subdomain1.xyz.com
- Stop incoming mail based on a deny list
- Stop relay based on the MAPS℠ external databases.

You can download the latest version of sendmail™ from http://www.sendmail.org.

MAPS Realtime Blackhole List (RBL℠)

The MAPS RBL℠ is a system for creating intentional network outages ("blackholes") for the purpose of limiting the transport of known-to-be-unwanted mass email. The MAPS RBL℠ is a subscription system, such that no one is ever denied connectivity to a non-RBL-subscriber. For more information about MAPS RBL℠, go to http://mail-abuse.org/rbl/.

To use the MAPS RBL, the following rules must be included in the sendmail™ configuration file, sendmail.cf:

```
R$*              $: $&{client_addr}
R::ffff:$-.$-.$-.$-   $: <?> $(host $4.$3.$2.$1.blackholes.mail-abuse.org. $: OK $)
R$-.$-.$-.$-      $: <?> $(host $4.$3.$2.$1.blackholes.mail-abuse.org. $: OK $)
R<?>OK            $: OKSOFAR
R<?>$+            $#error $@ 5.7.1 $:  Mail from $&{client_addr} rejected, see http://mail-abuse.org/rbl
```

MAPS Relay Spam Stopper (RSS℠)

The MAPS Relay Spam Stopper (RSS℠) is a query-able DNS-based database of spam-relaying mail servers. You can configure your mail server to utilize the RSS list. All emails (spam and non-spam) sent from the mail servers on this list will be rejected by your mail server. For more information about RSS℠, go to http://work-rss.mail-abuse.org/rss/index.html.

To use the MAPS RSS, the following rules must be included in the sendmail™ configuration file, sendmail.cf:

```
R$*              $: $&{client_addr}
R::ffff:$-.$-.$-.$-   $: <?> $(host $4.$3.$2.$1.relays.mail-abuse.org. $: OK $)
R$-.$-.$-.$-      $: <?> $(host $4.$3.$2.$1.relays.mail-abuse.org. $: OK $)
R<?>OK            $: OKSOFAR
R<?>$+            $#error $@ 5.7.1 $:  Mail from $&{client_addr} rejected, see http://mail-abuse.org/rss
```

MAPS Dial-up User List

The MAPS Dial-up User List (DUL℠) is a list of known dial-up or dynamically assigned pools of IP addresses. The list is available as a DNS zone arranged similar to the MAPS RBL℠, and also as a list of networks in CIDR format (Classless Inet-Domain Routing). You can configure your mail servers to use the DUL℠ to prevent users on those IP addresses from sending unsolicited email to your servers. For more information about the MAPS DUL℠, go to http://mail-

abuse.org/dul/.

To use the MAPS RSS, the following rules must be included in the sendmail™ configuration file, sendmail.cf:

```
R$*                $: $&{client_addr}
R::ffff:$-.$-.$-.$-    $: <?> $(host $4.$3.$2.$1.dialups.mail-abuse.org. $: OK $)
R$-.$-.$-.$-           $: <?> $(host $4.$3.$2.$1.dialups.mail-abuse.org. $: OK $)
R<?>OK              $: OKSOFAR
R<?>$+              $#error $@ 5.7.1 $:  Mail from $&{client_addr} rejected, see http://mail-abuse.org/dul
```

NOTE: Editing sendmail™ configuration file, sendmail.cf is rather complicated. You are recommended to use the M4 macro processor to generate new rules for sendmail™ configuration. This subject will be detailed next section.

Configure sendmail™:
There are configuration files that control how sendmail™ copes with mail from outside sources. Some are compiled database files, the others are simple text files. These are:

/etc/mail/sendmail.cf	The sendmail™ configuration file
/etc/mail/access.db	The "deny list". This is a database file of mail addresses from which to deny connections.
/etc/mail/aliases.db	Database for sendmail™ aliases
/etc/mail/relay-domains	List of domains from which relayed messages will be rejected. Messages from all other relays will be rejected.
/etc/mail/mailtertable.db	Database that maps host.domain names to special delivery agent and new domain name pairs.

Sample of the mailertable

```
subdomain1.xzy.com      esmtp:mt.subdomain1.xyz.com
subdomain2.xzy.com      esmtp:mt.subdomain2.xyz.com
subdomain3.xzy.com      esmtp:mt.subdomain3.xyz.com
```

Sample of the relay-domains file

```
subdomain1.xyz.com
subdomain2.xyz.com
subdomain3.xyz.com
```

sendmail™ configuration file
sendmail.cf is the main configuration file. It controls nearly everything. It's the most important file and is rather hard to edit manually. It is recommended to configure sendmail™ using M4 macro processor. Prior to sendmail™ version 8.10, M4 file was named *example*.m4. In version 8.10 and newer of sendmail™, M4 file name has been changed to *example*.mc. Tools in latest sendmail™ source expect *.mc* files. Below is a sample *.mc* file that incorporates local access

database (deny list), MAPS Realtime Blackhole List, MAPS Relay Spam Stopper and the MAPS Dial-up List into the existing sendmail™ configuration (sendmail.cf file).

```
divert(-1)
divert(0)dnl
VERSIONID(`example.mc 0.1')dnl
OSTYPE(`solaris2')dnl
DOMAIN(`generic')dnl
dnl
dnl Local SPAM filtering and mail routing
dnl
FEATURE(`access_db',`hash -o /etc/mail/access')dnl
FEATURE(`mailertable', `hash -o /etc/mail/mailertable')dnl
dnl
dnl MAPS definitions
dnl
FEATURE(dnsbl, `blackholes.mail-abuse.org',` Mail from $&{client_addr}
rejected, see http://mail-abuse.org/rbl')dnl
FEATURE(dnsbl, `relays.mail-abuse.org',` Mail from $&{client_addr}
rejected, see http://mail-abuse.org/rss')dnl
FEATURE(dnsbl, `dialups.mail-abuse.org',` Mail from $&{client_addr}
rejected, see http://mail-abuse.org/dul')dnl
dnl
dnl
FEATURE(`no_default_msa')dnl
FEATURE(`relay_hosts_only')dnl
FEATURE(`blacklist_recipients')dnl
define(`QUEUE_DIR', `/var/spool/mqueue/q*')dnl
define(`confSMTP_LOGIN_MSG', `$j *** FOR AUTHORIZED USE ONLY! ***')dnl
define(`confPRIVACY_FLAGS', `authwarnings noexpn novrfy needmailhelo')dnl
define(`confMCI_CACHE_SIZE',`5')dnl
define(`confTO_CONNECT',`5m')dnl
define(`confTO_COMMAND',`5m')dnl
define(`confTO_IDENT',`0')dnl
dnl
dnl Mailers
dnl
MAILER(local)dnl
MAILER(smtp)dnl
```

Note: The actual and complete .mc file should include other important configuration information such as maximum message size (MAX_MESSAGE_SIZE), maximum number of recipients per message (MAX_RCPTS_PER_MESSAGE), etc.

Let's name the above file example.mc. To create sendmail.cf from the example.mc file:

1) put example.mc file in /<path>/sendmail-x.xx/cf/cf
2) cd /<path>/sendmail-x.xx/cf/cf
3) ./Build example.cf (or make example.cf)
4) cp example.cf sendmail.cf

Rule Sets for Delivery by Sendmail™:

The rule sets for delivery are:

- Check local passwd file for local delivery

- Check Alias and forward database for delivery instructions
- Check the domain in the "rcpt to:" field. If that domain has a LUSER_RELAY, then forward the message to that system. If not, reject the message.

Detecting Spam

It is hard to detect spam by looking at the queue size in sendmail™, because sendmail™ processes the messages so quickly. Run a command like the one below on the mail servers and observe what happens. This example watches for mail for more than 30 addresses at once.

tail –f /var/log/syslog | grep to= | /usr/xpg4/bin/egrep –v 'nrcpts=[01234],' | grep nrcpts

IP Filter

You'll need to install IP Filter freeware to enhance spam filtering. You can download the latest version of this software from http://coombs.anu.edu.au/~avalon/ip-filter.html

IP Filter can

- Explicitly deny/permit any packet from passing through
- Distinguish between various interfaces
- Filter by IP networks or hosts
- Selectively filter any IP protocol
- Selectively filter fragmented IP packets
- Selectively filter packets with IP options
- Send back an ICPM error/TCP reset for blocked packets
- Keep packet state information for TCP, UDP and ICMP packet flows
- Keep fragmented state information for any IP packet, applying the same rule to all fragments
- Act as a Network Address Translation (NAT)
- Use redirection to setup true transparent proxy connections
- Provide packet header details to a user program for authentication
- In addition, supports temporary storage of pre-authenticated rules for passing packets through

This document does not include details on how to compile and install IP Filter. You can refer to IP Filter's How-To document at http://www.obfuscation.org/ipf/ipf-howto.txt for more details. IP Filter runs off a configuration file, ipf.conf, which is usually installed in /etc/opt/ipf/ directory. This configuration file contains all rules defined for IP packet filtering; and the rules are processed in order from top to bottom. You can edit this file manually. For example, in an effort to filter spam, you can set up these rules:

```
## block in-coming IP traffic on hme0 interface
block in on hme0 all head 200
block in log quick from 24.30.219.37/32 to any group 200      # spoof
block in log quick from any to 24.30.219.0/32 group 200       # smurf
block in  log quick from any to 24.30.219.63/32 group 200     # smurf
block in  log quick from 255.255.255.255 to any group 200     # from old-broadcast
block in  log quick from 127.0.0.0/8 to any group 200         # from localhost
```

```
block in  log quick from any to 127.0.0.0/8 group 200        # to localhost
block in  log quick from 10.0.0.0/8 to any group 200         # from unroutable
block in  log quick from 172.16.0.0/12 to any group 200      # from unroutable
block in  log quick from 192.168.0.0/16 to any group 200     # from unroutable
```

Conclusion

sendmail™ is a great tool to filter spam. ISPs can implement it to better ensure electronic communications for their subscribers. But this implementation cannot filter 100% junk mails. Many spammers continuously find ways to get around the blocks and filters. ISPs can only support anti-spamming up to a certain level. Therefore to better fight against spamming, the end users should also coordinate their effort with the ISPs and the government. When you get spammed, you can take legal action against the spammer if you want, but first you have to check to see there are any laws regarding this matter in your state. The best place to look for this information is at www.suespammers.org. And if you don't want to take legal action against the spammer, you can complain to the spammer's ISP. If the ISP doesn't take any action against the spammer, you may wish to file an RBL nomination against the IPS involved. You can find more about the RBL at http://mail-abuse.org/rbl/. Some other useful anti-spamming resources are:

- Fighting E-mail Spammers at http://eddie.cis.uoguelph.ca/~tburgess/local/spam.html
- The SPAM-L FAQ at http://www.claws-and-paws.com/spam-l/

References

1. "The Anti-Spam Home Page". URL: http://www.arachnoid.com/lutusp/antispam.html
2. "Anti-Spam Configuration Control". URL: http://www.sendmail.org/m4/anti-spam.html
3. "Unsolicited Commercial Email – The Problem". URL: http://www.cauce.org/about/problem.shtml
4. Reed, Darren. "IP Filter". URL: http://coombs.anu.edu.au/~avalon/ip-filter.html
5. Costales, Bryan *with* Allman, Eric. sendmail, Second Edition. O'Reilly and Associates, Inc., 1997.
6. "MAPS Realtime Blackhole List (RBL[sm])". July 06, 2001. Version 1.41. URL: http://mail-abuse.org/rbl/
7. MAPS Relay Spam Stopper (RSS[sm]). August 01, 2000. URL: http://work-rss.mail-abuse.org/rss/index.html
8. "MAPS Dial-up User List (DUL[sm])". July 18, 2001. Version 1.33. URL: http://mail-abuse.org/dul/
9. Mueller, Scott. "Promote Responsible Net Commerce: Help Stamp Out Spam! What is spam?". URL: http://spam.abuse.net/spam/whatisspam.html
10. Levine, John. "Promote Responsible Net Commerce: Help Stamp Out Spam! Why is spam bad?". URL: http://spam.abuse.net/spam/spambad.html