

ISPadmin
October, 2002
Stopping Spam, Part II

Introduction

In this installment of ISPadmin, techniques for stopping outbound spam (UCE originating on your network, destined for a machine on network for which you do not control) are examined. In the last edition, how to stop spam from the inbound side (from someone else's network to your mailbox) was covered in detail.

Background

Methods for stopping outbound mail are very different than those used to stop inbound spam. Most of the ways outbound spam are stopped can be classified as follows:

- Controlling access to a mail relay machine (for example, smtp.isp.net)
- Limiting SMTP access to known blocks of open mail relays (for example, Korea)
- Limiting the number of outbound SMTP connections a client can make over a period of time
- Capping the amount of k/sec an outbound SMTP connection can make

The methods covered in this article will fall into one of the categories listed above, although the coverage will be grouped differently to enable clearer coverage of the topics.

Generic Methods

First, lets discuss generic methods which are not tied directly to a specific open source solution or network hardware (for example, routers). These methods can be applied to any mail infrastructure, though sendmail specific information is listed within this section.

==>Restricting IP

Controlling what IP address are allowed to send mail through a mail server is an important step everyone who runs a mail system on the Internet should take. This is a very common method to control access to a mail relay. In the provider's mail relay machines, a list of IP addresses or blocks is kept that are allowed to relay mail through the relay(s). For sendmail, the "IP allowed to relay" list is kept in an access database entry similar to the following:

```
209.206.10 RELAY
```

(Sendmail access databases were covered in last issue's ISPadmin column.) Even if you are not a provider, if you are running sendmail you should be restricting access to your mail relays in this manner. If you don't, you run the very high risk of becoming a spam pariah!

==>POP before SMTP

The POP before SMTP method requires the end subscriber to simply check their mail before sending it. This method can be used for "roaming" subscribers, who won't be coming from one of the provider's own IP address ranges. Once the POP box is accessed successfully, the subscriber's IP address goes into the IP address "allowed" list on the mail relay(s) for a certain period of time, most commonly 30 minutes. In the case of a sendmail based mail relay, the method to control mail relay access can be performed via the access database entry, identical to the approach outlined in the "Restricting IP" section.

==>Mail Message Metering

(Disclaimer: This author developed the Mail Message Metering anti-spam method, which has a patent pending. Describing the method here does not imply the ability to use the system described here.) The Mail Message Metering method is simple in concept, and relatively simple to implement. The method is useful to wholesale Internet access providers, although any enterprise that generates lots of outbound mail could use it.

As each subscriber generates an outbound mail message, the network component (switch, RAS gear, DSL aggregating equipment, etc.) redirects the connection to a specially configured mail relay. This specialized mail relay queries a database which contains a current listing of all originating IP addresses that have relayed mail, and associated counts of the number of messages for several time periods (for example, past minute, past 30 minutes and past hour). If the message would exceed predetermined thresholds, then the message would be re-queued. If the message didn't exceed the limits, then the message would be allowed through and the counts updated appropriately.

Other people and organizations hold anti-spam patents. Of these, Brightmail is probably the best known. However, this author (who is not an attorney) can find no patent (granted or pending) specific to outbound spam.

The benefits of this approach are many:

- Blocks high percentage of outbound spam
- No subscriber and little customer impact
- Configurable and scalable
- Limited impact on authentication (RADIUS) servers

The shortcomings are:

- Requires "white hat" list of legitimate bulk mailers
- Requires use of SMTP redirection (may require additional hardware)

The December, 2000 issue of ;login: contained an in depth article on the Mail Message Metering solution.

Open Source Packages

One open source packages is specifically designed to counter outbound spam (kai's SpamShield). The others others described can be used to control both inbound and outbound spam.

==>kai's SpamShield 1.0

kai's SpamShield is probably one of the oldest packages out there specifically designed to counter outbound spam. It is a perl script run out of cron which works by analyzing the most recent sections of the sendmail log file (usually maillog). The program counts the IP addresses from which messages are originating. If these counts exceed previously entered thresholds, the sender's access to the mail relay is blocked. While it is dated (it doesn't appear to have been updated since 1997) it is very effective against outbound spam.

kai's SpamShield version 2.0 was just announced as of this writing in July, 2002. No details on the functionality included in the new version exist on the web site, however.

==>Blackmail

Blackmail performs various checks against the headers of incoming and outgoing mail messages. These checks include:

- Known sources of spam
- Specific words and/or phrases
- Resolvable names in headers
- Black hole lists
- To: and From: headers
- Correct header formation

While more recent than kai's SpamShield, it appears that most of these checks are performed by spamassassin as well. One difference would be the fact that Blackmail is written in C, while spamassassin is written in perl.

==>Procmal

System wide procmal filters can be built to assist in the fight against spam. Two such packages are "The SpamBouncer" and "Email Sanitizer". These work by encapsulating the various anti-spam rulesets (for example, black hole lookups, resolvable to/from domains, etc.) as procmal recipes. While this author has no direct experience with them, there are enough procmal based tools out there to indicate this is a valid approach.

==>SMTP proxy

SMTP proxies (such as Obtuse Systems Corporation's Juniper firewall toolkit or Trusted Information Systems fwtk) contain basic SMTP filtering that can be used to control outbound spam. In fact, the Mail Message Metering implementation utilized the Juniper firewall toolkit's smtpd as the basis for the message processing. The proxy approach is a minimalistic one, as spamassassin contains much more anti-spam functionality built into it. However, they are implemented in C/C++ which may make the proxies more reliable than code written in perl.

Stopping spam at the Network

There are ways spam can be controlled by the provider at the network level:

- Blocking access to known open relays via access control lists (ACL's) on routers
- Caller-ID blocking

The downside to these methods is they do take resources on the network components (such as routers), which can cause additional cash outlays by the provider to implement these methods.

==>Blocking access to known open relays

One very effective (but drastic) way to reduce unwanted outbound spam is to simply disallow access to all SMTP servers except for the provider's own mail relays. This could be accomplished by the following ACL on a Cisco router:

```
access-list 101 permit tcp host a.b.c.d any eq smtp
access-list 101 permit tcp host e.f.g.h any eq smtp
.
.
.
access-list 101 deny tcp i.j.k.l.0 0.0.0.255 any eq smtp
access-list 101 permit ip any any
```

The first two access-list statements allow access to legitimate mail relays, and more permit hosts/networks could be added. The third access-list statement denies access all other access to port 25 (SMTP) outside what is specified in the permit list. The final statement allows all other traffic to be routed normally.

A variation on this idea is to block outbound SMTP access to known networks that house open relays, such as Korean networks. A dial up customer should be using the mail relays provided, rather than misconfigured ones located halfway around the world!

==>Other RAS/network techniques

Many spammers will block caller ID to make it harder to track the abusers down. One technique that is used to block spammers from wholesale dial up networks is to disallow outbound SMTP access to anyone who calls in without providing caller ID. This will stop a lot of spam. Also, RAS filters can be loaded dynamically on to the modem ports via RADIUS, allowing SMTP access to a certain set of IP addresses, and excluding the rest. In fact, UUNET *requires* its customers to pass a RADIUS attribute (Ascend-Data-Filter), allowing outbound SMTP access to its wholesale customers' mail relay, and nothing else.

Other tactics that can be tried (with additional network hardware) might be to limit the outbound SMTP connection rate, or outbound SMTP bandwidth, coming from a particular IP address. This author is not aware where this has been tried "in the wild" on a production network.

Miscellaneous Topics

This section contains odds and ends regarding both inbound and outbound spam.

==>Acceptable Use Policy

Perhaps the most important document a service provider has is its Acceptable Use Policy or AUP. Without a properly written AUP, it is impossible to legally shut off customers who abuse a providers network. All organizations, be they providers, small companies, large companies, non profits or others should have an AUP. While it takes time and effort to write a good one, the headache it solves in the long run is well worth it.

==>Legal aspects

A book could be written on the legal aspects of UCE. In the US at this point in time, the only laws governing spam at the federal level surround fax broadcasting (governed by the Federal Communications Commission), and the legality of claims made by spammers (governed by the Federal Trade Commission). Case law is being built every day. In July, 2002, Earthlink was awarded US\$25 million in a lawsuit against spammers. The FTC has been active in pursuing spammers who make illegal claims.

In the US, the only codified anti-spam law is at the state level. David E. Sorkin has a great site that summarizes current status of anti-spam law, both inside and outside the US.

==>Staff

At most ISP's, customer support and/or the network operations center personnel handle spam complaints. At Ziplit, the company dedicated approximately two staff positions to handle the influx of spam complaints, with a 70,000 port dial in network. Many complaints are duplicates, or are sent in error, which causes additional overhead.

Automated systems such as Spamcop work well. However, they are not infallible and do make mistakes. One benefit of such systems is the elimination of duplication of effort automated systems can provide. Spamcop will stop sending spam reports to the provider, once the provider tells Spamcop the spammer has been deactivated. However, Spamcop continues to send duplicate spam reports, with the same "footprint" (ie, source IP address, subject line, etc.) until the provider takes action.

==>Costs

The additional strain spam puts on staff, machines and networks is hard to quantify. If we use an assumption that 33% of all email is spam, that loosely translates into 33% higher costs for the provider. The two additional staff positions could be eliminated, if spam was not a problem. A server or two could probably be reallocated at a small to mid size ISP, while a larger provider could probably eliminate more. The upstream network connections, if the provider buys transit, would be less without spam.

==>Usenet News spam

Most news servers these days are able to control news spammers without much difficulty. InterNetNews (INN) v2.3.2 has an "exponential backoff" feature. The associated control parameters are:

backoffauth
backoffdb
backoffk
backoffpostfast
backoffpostslow
backofftrigger

Check the man pages for inn.conf and search for "backoff" for more information. If the Highwinds Software series (typhoon/cyclone/twister) of news servers is used, a perl program is available to rate limit article posting. This rate limiting works very well.

==>Places to send your spam

Ever wonder where you can send spam you receive (besides to the provider that originates it)? A list of email addresses appears below; if anyone knows of additional email addresses to send junk mail to, please send them and they will be published in a future column. Some of these addresses are just statistics trackers, others are for actual complaints, and some are commercial services that block spam and who use the email to generate rules for protecting their customers. Here are some email addresses this author is aware of:

spamrecycle@chooseyourmail.com	The spam recycling center (statistics)
uce@ftc.gov	FTC's junk mail address
fraud@uspis.gov	For complaints involving US Postal mail

enforcement@sec.gov	For securities related complaints involving US
publicly listed companies	
cyberfraud@nasaa.org	For securities related complaints involving US
publicly listed companies	
otcfraud@cder.fda.gov	For food/drug related complaints
junk@brightmail.com	Honeypot address for Brightmail spam filtering
service	

Conclusion

There are available tools for ISP's (and others) to control outbound spam. Mail transfer agents (MTA's) such as Sendmail can be configured to allow certain IP address ranges to relay mail, which all organizations running a mail server on the Internet today should employ. Outside of MTA's, kai's SpamShield can be utilized to control outbound spam, and other mail proxy agents can be useful as well. These open source methods work, but are not perfect and take effort to implement. Steps can be taken at the router/network device level as well, but these are not adaptive and must often be regularly updated. Some proprietary methods (such as Mail Message Metering) do exist, but are applicable to certain classes of spam sources (such as large ISP's) and covered by intellectual property law.

References

POP-before-SMTP: <http://popbsmtp.sourceforge.net/>
 Relay control in sendmail for roaming users:
<http://www.sendmail.org/~ca/email/roaming.html>
 Mail Message Metering: <http://www.ziplink.net/ziplink/solutions/mmm/>
 Kai's SpamShield: <http://spamshield.conti.nu/>
 Obtuse Systems Juniper firewall toolkit smtpd: <http://www.obtuse.com/smtpd.html>
 Blackmail: <http://www.jsm-net.demon.co.uk/blackmail/blackmail.html>
 The SpamBouncer: <http://www.spambouncer.org/>
 Email Sanitizer: <http://www.impsec.org/email-tools/procmail-security.html>
 TIS fwtk: <http://www.fwtk.org/fwtk/>
 ISP-Planet article on Earthlink spam lawsuit: <http://www.internetnews.com/isp-news/article.php/1430591>
 David E. Sorkin's spam law site: <http://www.spamlaws.com/>
 Spamcop: <http://spamcop.net/>
 INN: <http://www.isc.org/products/INN/inn-current.html>
 Highwinds Software (Typhoon/Cyclone/Twister): <http://www.highwinds-software.com/discussion/index.html>
 SpamCon Foundation's list of places to send junk email:
<http://www.spamcon.org/recipients/spam-response/help-statistics.shtml>
 Brightmail: <http://www.brightmail.com/>
 SpamCon Foundation: <http://www.spamcon.org/>
 Brightmail anti-spam patents: <http://patft.uspto.gov/netacgi/nph-Parser?Sect1=PTO2&Sect2=HITOFF&u=%2Fnetacgi/nph->

adv.htm&r=0&p=1&f=S&l=50&Query=in%2F%22paul%3B+sunil%22%0D%0A&d=ft
00