Spam Control

Devdas Bhagat
devdas.b@gmail.com

Definitions

Spam is Unsolicited Bulk Email

SPAM is canned pig meat.

Devdas Bhagat
devdas.b@gmail.com

The spam cartel

- Who are the major players?
  - People sending spam
  - Major companies who send spam
  - ISPs with an eye to make a quick buck
  - Data center operators
  - Malware developers
  - Organised crime
  - Domain Registrars
  - Home users with inadequate security

Devdas Bhagat
devdas.b@gmail.com

The spam cartel (contd)

- People sending spam
  - Some people make a living off spam
  - A very good living, in fact
    - http://news.bbc.co.uk/2/hi/business/3581435.stm
  - Such spam is mostly for selling products
    - Drugs
    - University diplomas
    - Sexual enhancement products

Devdas Bhagat
devdas.b@gmail.com

The spam cartel (contd)

- Mainline companies sending spam
- Also termed as mainsleaze
  - This is often stupid marketers or marketing departments
- Alternatively, there are cultural issues involved
  - Some cultures think it right to send out an ad in email, because it is “informative”.
  - This is a culture clash on the Internet

Devdas Bhagat
devdas.b@gmail.com

The spam cartel (contd)

- ISPs out to make a quick buck
  - Some ISPs see benefits in hosting spammers who pay a lot of money for connectivity
  - Also advertised as bulletproof hosting
- Datacenter operators who refuse to terminate spammy customers
- Also ISPs who refuse to terminate botnet command and control systems
  - See Intercage/Atrivo

Devdas Bhagat
devdas.b@gmail.com

The spam cartel (contd)

- Domain name registrars
  - These people share a love-hate relationship with spammers.
  - Spammers run through a lot of domain names, and hence are large volume customers in the domain name business.
  - However, they negatively impact the reputation of the registrar from whom they buy domains

Devdas Bhagat
devdas.b@gmail.com

The spam cartel (contd)

- Malware developers
  - People are paid to write undetectable bots
  - Botnets are rented out to spammers
    - For a few hundred dollars per spam run
    - Botnets are also used for DDoS attacks.
    - And to host illegal content

Devdas Bhagat
devdas.b@gmail.com

The spam cartel (contd)

- Organised crime
  - Fraud (419 type scams)
    - This is big business.
      - Literally in the 100s of millions
  - People sending trojans or keyloggers
  - Identity theft
    - Another big business
      - Estimates say 100 Billion USD
  - Pirated software downloads
  - Links to websites installing malware

Devdas Bhagat
devdas.b@gmail.com

The spam cartel (contd)

- Users with inadequate security
- We usually treat these people as victims
- They wittingly or unwittingly participate in the spam ecosystem
- The situation has been made worse with the rise of always connected broadband systems.

Devdas Bhagat
devdas.b@gmail.com

Controlling spam

Devdas Bhagat
devdas.b@gmail.com

This is a hard problem

- Spam is a social issue, not a technical problem
- Controlling social problems is hard
  - Humans suck
- You need a complex, multi-pronged approach to solve this problem.
  - Technical
  - Legal
  - Social Norms

Devdas Bhagat
devdas.b@gmail.com

The legal solutions

- Various countries have passed antispam laws
- Or at least, laws which pretend to be anti-spam
- Good examples
  - Australia
  - Europe
- Bad examples
  - The US CAN-SPAM act

Devdas Bhagat
devdas.b@gmail.com

Legal solutions aren't enough

- Enforcing the law takes time
- You still need to catch the criminals
- Spam is often cross-border crime
  - Extradition issues come into play
- Given that this is a criminal offense, who can sue may be restricted
- The financial losses involved are often small, so it may not be worth filing a lawsuit

Devdas Bhagat
devdas.b@gmail.com

Social norms

- Changing social norms so that spam is seen as unacceptable behaviour is the best solution
- However, social change is slow
- This needs a concerted effort on educating people about spam, the consequences and getting them to change their behaviour

Devdas Bhagat
devdas.b@gmail.com

Possible changes

- Don't buy from spammers
  - Boulder pledge
- Convey the impression that spamming is rude
  - This is a direct attack on people assuming that free advertising over email is the same as that paid for when buying web ads.
- Educate end-users on network security best practices
  - http://www.ranum.com/security/computer_security/editorials/point-counterpoint/users.html

Devdas Bhagat
devdas.b@gmail.com

Technical options

NONE OF THESE ARE REALLY EFFECTIVE

Devdas Bhagat
devdas.b@gmail.com

Outbound spam

This is probably where you need to put in most of your efforts.

Devdas Bhagat
devdas.b@gmail.com

At the edge

- Block port 25/tcp outbound from all dynamic IP blocks on your network.
  - This is a brutally efficient solution to outbound spam.
- Require your customers to do the same things on their networks.
  - This may or may not be feasible, but it can be made part of your customer agreement and AUP.

Devdas Bhagat
devdas.b@gmail.com

Linux Router

```
# Allow traffic to the submission port
/sbin/iptables -I FORWARD -s dynamic/block-p tcp --dport 587-j ACCEPT

# Reject traffic from dynamically allocated addresses to mailservers
# Port 25 is for MTA to MTA traffic
/sbin/iptables -I FORWARD -s dynamic/block -p tcp --dport 25 -j REJECT
```

Devdas Bhagat
devdas.b@gmail.com

Cisco outbound acl

```
interface ethernet0
ip access-group 101 out
!
access-list 101 permit host 192.0.20.1 any eq smtp
access-list 101 deny any any eq smtp
```

Devdas Bhagat
devdas.b@gmail.com

Controlling spam at the MTA

- Separate your MX, submission and outbound servers.
  - Do not try and optimise this onto one box, or one MTA instance.
- Accept end user email on port 587/tcp
  - Require authentication, and preferably TLS
  - SMTP AUTH is a standard
  - Every MUA supports it, or supports an alternative mechanism for message submission

Devdas Bhagat
devdas.b@gmail.com

Things you can't use

- DNSBLs
  - These are meant to block email at MX hosts, not submission hosts
  - There are multiple instances of providers trying to use DNSBLs against their own authenticated users
  - This is a silly mistake when trying to prevent spam relaying

Devdas Bhagat
devdas.b@gmail.com

What works

- Rate limiting message sending by authenticating user and IP address
- Content filtering/content analysis
  - This is difficult to get right
  - This is CPU intensive
  - But done right, it's brilliant at cleaning up your outbound stream.

Devdas Bhagat
devdas.b@gmail.com

Content filtering via simple checks

- You can try and match simplistic header/body patterns
  - Things like S. 1618 in the body of the message
- These aren't sufficiently granular
- You will end up getting false positives
- Some levels of coarse grained control is useful though

Devdas Bhagat
devdas.b@gmail.com

Spamassassin

- You can inject SA into your outbound mail pipeline
- SA runs a bunch of tests
  - It applies a fairly wide range of heuristics
- Remember to disable DSNBLs
- Train the Bayesian filter
  - This is hard manual labour
  - It needs to be done once, and then just occasionally checked for accuracy

Devdas Bhagat
devdas.b@gmail.com

amavisd

- Amavisd is a wrapper around SA and one or more antiviruses
- This is reasonably efficient at per user preferences and access control
  - It supports a variety of backends and policy statements

Devdas Bhagat
devdas.b@gmail.com

INBOUND SPAM

Devdas Bhagat
devdas.b@gmail.com

Controlling inbound spam

- This is a mostly solved problem
  - Well, not really
  - We actually have a fairly good handle on controlling spam directed at our networks, and can easily stop over 99% of spam upfront
  - The bit users actually complain about is the last one percent
    - But one percent of a big number is still a big number
- Defence in depth is a standard strategy

Devdas Bhagat
devdas.b@gmail.com

On the router

BGP or routing ACLs from Team Cymru

http://www.team-cymru.org/Services/Bogons/

Devdas Bhagat
devdas.b@gmail.com

A quick note on SMTP

- SMTP is designed to facilitate communication
- SMTP is a synchronous, request-response protocol
- Every SMTP command verb expects a response.
- SMTP clients MUST NOT send a new command unless PIPELINING is allowed.
  - In which case, they may send commands and then wait for responses until DATA.

Devdas Bhagat
devdas.b@gmail.com

Tactics

- Accept email only for valid recipients
- Local black lists
- DNSBLs
- SMTP client data sending behaviour
  - Delaying the initial greeting
  - Greylisting
- SMTP client network behaviour

Devdas Bhagat
devdas.b@gmail.com

Tactics

- Looking for invalid content
- Heuristics
  - Malware detection engines
  - Content heuristics via pattern matching
  - Bayesian analysis

Devdas Bhagat
devdas.b@gmail.com

On the MX

- Recipient validation
- Postfix
  - reject_unauth_destination
  - reject_unlisted_recipient

Devdas Bhagat
devdas.b@gmail.com

DNSBLs

- These are DNS based checks for IPs matching a given condition
  - This may be IP addresses found to originate spam, dynamically assigned IP address blocks, or anything else.
- Check the policy of listing before using a DNSBL
- Good options: Spamhaus has a few DNSBLs.

Devdas Bhagat
devdas.b@gmail.com

Building your own DNSBL

- Use RBLDNSD.
- http://www.corpit.ru/mjt/rbldnsd/
- Do NOT use BIND. BIND will suck up all your memory, and cause your machines to misbehave.

Devdas Bhagat
devdas.b@gmail.com

rbldnsd config

```
# If the first character is a :, then that line is the response message.
# $ is replaced by the IP address.
:127.0.0.2:$ has been blacklisted. See http://www.example.com/search?q=$ for details
192.168
10.0.0.1-15
192.0.20.0/24
```

Devdas Bhagat
devdas.b@gmail.com

SMTP client behaviour

- Spam clients are usually not written to RFC compliance
- Requiring RFC compliance leads to a slight reduction in spam
- Greylisting works by first temporarily rejecting a message, and then allowing it when the same (client IP, sender, recipient) triplet is retried.
- Greet pause delay delays the initial 220 greeting. This causes some spam engines to break.

Devdas Bhagat
devdas.b@gmail.com

SMTP client network behaviour

- Use the BSD packet filter pf, which has passive OS detection built in.
- Linux has a third party iptables module which does this
  - See http://www.ioremap.net/projects/osf
- There are a number of operating systems which have no business connecting directly to a MX on port 25
  - Like Windows XP.

Devdas Bhagat
devdas.b@gmail.com

More complex checks

- Examining headers
- Simple pattern matching
- Milters

Devdas Bhagat
devdas.b@gmail.com

Example configurations (outbound)

```
content_filter = smtp-amavis:[127.0.0.1]:10024
milter_default_action = accept
smtpd_helo_required = yes
smtpd_milters = inet:localhost:7357
smtpd_recipient_restrictions = permit_mynetworks
        check_policy_service inet:127.0.0.1:10032
        permit_sasl_authenticated
        reject_unauth_destination
smtpd_sasl_auth_enable = yes
smtpd_sasl_authenticated_header = yes
smtpd_sasl_local_domain =
smtpd_sasl_security_options = noanonymous
strict_rfc821_envelopes = yes
```

Devdas Bhagat
devdas.b@gmail.com

Example configurations (inbound)

```
header_checks = regexp:/etc/postfix/header.re
mime_header_checks = pcre:/etc/postfix/mime_header.pcre
relay_domains = hash:/etc/postfix/relay_domains_map
relay_recipient_maps = regexp:/etc/postfix/relay_receipt_map
smtpd_recipient_restrictions = permit_mynetworks
        reject_unauth_destination
        reject_invalid_helo_hostname
        check_helo_access hash:/etc/postfix/helos
        check_helo_access regexp:/etc/postfix/helo_regexp
        check_client_access hash:/etc/postfix/whitelisted_clients
        check_client_access hash:/etc/postfix/clients_access
        check_sender_access pcre:/etc/postfix/sender.pcre
        check_sender_access     hash:/etc/postfix/sender_access
        check_recipient_access  hash:/etc/postfix/recipient_access
        check_helo_access proxy:pgsql:/etc/postfix/postgres-helo.cf
        check_recipient_access proxy:pgsql:/etc/postfix/postgres-check_recipient_maps.cf
        reject_unlisted_recipient
        reject_unknown_reverse_client_hostname
        reject_rbl_client       zen.spamhaus.org
        warn_if_reject reject_rbl_client       dnsbl.sorbs.net
        check_policy_service inet:127.0.0.1:10031
```

Devdas Bhagat
devdas.b@gmail.com

```
$cat /etc/postfix/header.re
# This file needs to be maintained carefully.
# The data here is an example, please use your own.
/^Thread-Index: Aca6Q/          REJECT          Spam spoor.
/^Received: from.* \(Postfix\) with ESMTP id \S+\s*$/   REJECT    Missing TZ

$cat /etc/postfix/mime_header.pcre
# Slightly modified from Jim Seymour's page:
# http://jimsun.linxnet.com/misc/header_checks.txt
/^Content-(Disposition|Type):\s+.+?(?:file)?name="?.+?\.(386|ad[ept]|app|as[dpx]|
ba[st]|bin|btm|cab|cb[lt]|cgi|chm|cil|cla(ss)?|cmd|cp[el]|crt|cs[ch]|cvp|dll|dot|drv|
ex[_e]|fon|fxp|hlp|ht[ar]|in[fips]|isp|keyreg|ksh|lib|lnk|md[abetw]|mht(m|ml)?|ms[cipt]|
nte|nws|obj|ocx|ops|ov.|pcd|pgm|pif|prg|sc[rt]|sh[bs]?|slb|smm|sw[t]|sys|vb[esx]?|
vir|vmx|vxd|wm[dsz]|ws[cfh]|xms|\{[\da-f]{8}(?:-[\da-f]{4}){3}-[\da-f]{12}\})\b/
    REJECT ".$2" file attachment types not allowed. Please zip and resend.
/^Content-(Disposition|Type):\s+.+?(file)?name="?.+?\.com(\.\S{2,4})?(\?=)?"?(;|$)/
    REJECT ".com" file attachment types not allowed. Please zip and resend.
```

Devdas Bhagat
devdas.b@gmail.com

```
$ cat /etc/postfix/helos
# No one should ever HELO/EHLO with these names. These
# are usually misconfigured desktops.
localhost                    550          You aren't localhost.
localhost.localdomain        550          You aren't localhost.localdomain.

$ cat /etc/postfix/whitelisted_clients
72.14.246             OK
67.15.184             OK
67.15.47             OK
64.34.209.213          OK
64.34.200.165          OK

$ cat /etc/postfix/sender_access
tizabal@icqmail.com REJECT
nosy.biz REJECT
sigbandzlawdyj@bandzlaw.com REJECT
nogbaandersfiw@baanders.com REJECT
fidchat@onelist.com REJECT
```

Devdas Bhagat
devdas.b@gmail.com

Policy daemons

- A policy daemon is an external service which Postfix will ask to make decisions.
- The decision can be anything which Postfix will accept in normal configuration.
- http://www.postfix.org/SMTPD_POLICY_README.html
- Writing a policy server is **easy**. Use your favorite scripting language.

Devdas Bhagat
devdas.b@gmail.com

Abuse issues

- Most large providers run feedback services, termed as FBLs
- FBL send spam complaints along with data to a specified address in ARF – The Abuse Reporting Format
- ARF is designed for easy processing to get the spam out of the complaint

Devdas Bhagat
devdas.b@gmail.com

SPF, DKIM, etc

- These aren't anti-spam solutions.
- SPF tries to give the owner of the domain control over where email claiming to be from that domain will originate.
  - This has problems with email forwarding
- DKIM is simply signing SMTP message headers by the SMTP server to verify authenticity.
  - DKIM signed headers with a valid signature are not spoofed and can be trusted.

Devdas Bhagat
devdas.b@gmail.com

Stuff which doesn't work

- Captchas or equivalent
  - Just promise humans free pictures or money
- Trying to charge per message received
  - This breaks legitimate mailing lists
- Callbacks
  - Challenge Authentication Response Protocol
- Hashcash
  - Spend CPU time to allow email to go through

Devdas Bhagat
devdas.b@gmail.com

Additional resources

- MAAWG – The Messaging Anti-Abuse Working Group
  - http://www.maawg.org/ is the industry association dealing with abuse.
  - Their website has a lot of useful implementable policy suggestions against malicious activities.

Devdas Bhagat
devdas.b@gmail.com

To summarise

- Spam is a social problem
- Using technology to stop social problems doesn't work
  - That's what lawyers are for
- Spam isn't going away, just like cockroaches
- Spam can be migitated

Devdas Bhagat
devdas.b@gmail.com

Devdas Bhagat
devdas.b@gmail.com