Extrusion Detection

Richard Clayton

Security Group Seminar
Cambridge

3rd February 2004

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

thus™

Some jargon

• Spam	unsolicited bulk email (UBE, UCE etc)
• Smarthost	email server used as sole destination
• Bounce	to reject email with a failure report
• Relay	to send via an intermediate machine
• Proxy	to provide a service on behalf of others
• Virus	self-replicating malware (aka Worm)
• Extrusion	a bad pun on “intrusion” detection

Summary

- A short history of “spam”
- Logging by email “smarthosts”
- Spotting open servers
- Spotting viruses
- Mail loops
- False positives & clueless behaviour
- Where next ?

NB: all examples are anonymised!

A short history of spam

- 1994: Individual senders
- 1995: Throwaway accounts
- 1996: Open relays
- 1997: Dedicated systems **`friend@public.com`**
- 1998: Proxy hijacking **`wingate`**
- 2001: Brute forcing SMTP AUTH
- 2003: Trojan networks

All dates are approximate

Current (Feb 04) problems for ISPs

- SMTP AUTH
  - Exchange “admin” accounts + *many others*
- Open proxies
  - mainly “trojans on non-standard ports”
- Systems still insecure “out of the box”
  - brand new XP is compromised before secured
- Blacklisting of IP ranges & smarthosts
  - **`listme@listme.dsbl.org`**

more people get on with us

Where do lists come from ?

- Usenet
- Archives
- FAQs and other articles
- Web sites
- Mailing lists
- Magazine subscriptions & trade shows
- Imagination
- OTHER LISTS!

more people get on with us

Who are they sending it to ?

• Me ! <richard@turnpike.com>	14 last week
other addresses	6 last week
• Abuse desk	
<abuse@demon.net>	100 emails in September
• Robots	
<faq_request@turnpike.com>	25% of 1998 activity
• Invented names	
<merchantlike@highwayman.com>	20 a week
• Dead names	
<richard@locomotive.com>	22 last week
last used in July '96	

Guaranteed "spam" in 2003

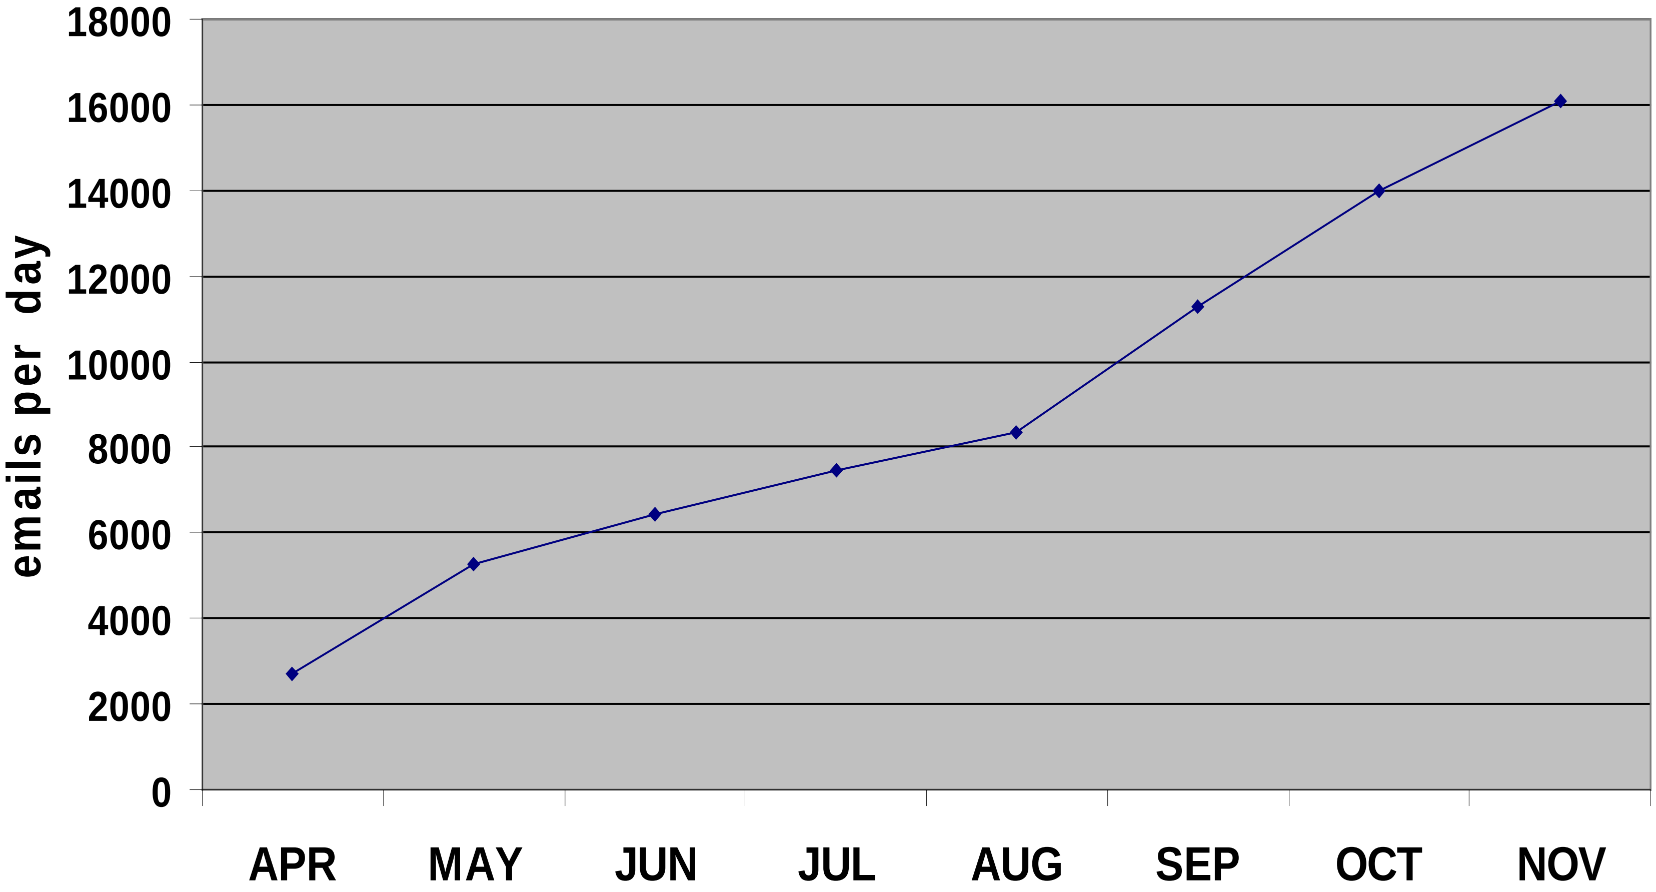

If I am rejecting this amount, why haven’t I been seen as a problem?

Smarthost logging summaries

- Traditional reports give top 50 sources and destinations
  - intended for load balancing
  - used to spot mail loops
  - used to spot outrageous abuse
- They are often irrelevant (hence unread)
  - sysadmin primary concern is service protection, not with more general abuse issues

Smarthost logging details I

These are for Exim v3 (YMMV)

• IP address	*for Demon this is account*
• HELO	*identity of source*
• MAIL FROM	*the “envelope” sender*
• Size	*number of bytes*
• Message ID	*supposedly unique*

Smarthost logging details II

Continued...

- RCPT TO *actual destination*
  - successful delivery (or to “fallback”)
  - failures (5xx)
  - necessity to retry (no response or 4xx)
- Other reports
  - restarts, system failures etc

NO CONTENT IS RECORDED

2004-02-02 09:55:54

09:55:54

```
1Anao3-0004b6-0U <= Accounts@example.co.uk H=mailgate.example.co.uk
               (basil.example.co.uk) [212.240.1.2] P=esmtp S=47563 id=XXXX@BASIL
```

09:55:55

```
1Anao6-0004Pi-0U <= FILTER-DAEMON@example2.co.uk H=example.demon.co.uk (e800-
               lam1.sbs.local) [193.237.1.1] P=esmtp S=1485
               id=FSHWTDMFZV.MOX.quarantine@example2.co.uk
1Anao6-0004bm-0U => nichola@example1.com R=lookuphost T=remote_smtp
               H=mail2.example1.com [212.78.2.3] C="250 2.0.0 i129ttO01977 Message
               accepted for delivery"
1Anao6-0004bm-0U Completed
1Anano-0004Vn-0U ** fwDOpQDeUf@puebla.com R=error_message T=remote_smtp: SMTP error
               from remote mailer after RCPT TO:<fwDOpQDeUf@puebla.com>: host
               puebla-com.mr.outblaze.com [205.158.62.147]: 550 1Anano-0004Vn-0U
               fwDOpQDeUf@puebla.com: error ignored
1AnYor-000KPc-0U SMTP error from remote mailer after RCPT TO:<handtech@email.com>:
               host email-com.mr.outblaze.com [205.158.62.23]: 450
               <handtechhandtech@YAHOO.COM.JP>: No thank you: rejected: Domain not
               found
1Anao4-0004Zn-0U => slotcarshop@example.net R=lookuphost T=remote_smtp
               H=mail.example.net [196.25.0.1] C="250 ok 1075715185 qp 22534”
```

09:55:58

```
1Anao3-0004b6-0U => yo_takenaka@example.com R=lookuphost T=remote_smtp
               H=mailsweeper.example.com [194.130.1.2] C="250 192.168.0.12: Message
               accepted for delivery"
1Anao3-0004b6-0U Completed
```

Simplification

- Process logs into single line per email
  - discard logs for system machines
  - read all details (discards info at midnight)
  - record failure status
    - remote site may indicate is spam
  - encode difficult characters (commas etc)
  - sort the resultant 300MB+ file
- About 1200 lines of (professional) Perl

[062.049.001.001] 2004-01-18

```
16:03:16 1AiFOL-000LKL-0V 0,4650,rod@example.com,example.demon.co.uk,[192.168.0.1],
         BC305E52.17B7%%rod@example.com,<accounts@example.com>
16:08:09 1AiFT6-000Lnb-0V 0,1262,rod@example.com,example.demon.co.uk,[192.168.0.1],
         BC305F7E.17B8%%rod@example.com,<media@example.com>
16:34:27 1AiFsY-000DVK-0Z 0,4305,rod@example.com,example.demon.co.uk,[192.168.0.1],
         BC3065A8.17BD%%rod@example.com,<Jessy.Hunter@example.co.uk>
16:35:28 1AiFtX-000Dcs-0Z 0,995,rod@example.com,example.demon.co.uk,[192.168.0.1],
         BC3065E5.17BE%%rod@example.com,<jules@example1.co.uk><adgewiseman@example2.co.uk>
16:36:40 1AiFuc-000DkB-0Z 0,102305,rod@example.com,example.demon.co.uk,[192.168.0.1],
         BC306628.17BF%%rod@example.com,<media@example.com>
16:40:02 1AiFxx-000EBV-0Z 0,1672,rod@example.com,example.demon.co.uk,[192.168.0.1],
         BC3066F7.17C1%%rod@example.com,<!vik_lok@yahoo.com>
16:40:02 1AiFxy-000EC8-0Z 0,2712,<>,<bounce>,1AiFxx-000EBV-0Z,,<rod@example.com>
16:44:41 1AiG2T-000EjP-0Z 0,4174,rod@example.com,example.demon.co.uk,[192.168.0.1],
         BC306810.17C3%%rod@example.com,<jeff@example2.com>
18:24:34 1AiHaj-000Jzn-0W 0,480842,rod@example.com,example.demon.co.uk,[192.168.0.1],
         BC307F5F.17CB%%rod@example.com,<media@example.com>
18:29:34 1AiHfx-000Kyi-0W 0,3263,rod@example.com,example.demon.co.uk,[192.168.0.1],
         BC3080A5.17CD%%rod@example.com,<specialprojects@lexample.net>
18:31:36 1AiHhw-000LSf-0W 0,3178,rod@example.com,example.demon.co.uk,[192.168.0.1],
         BC30811F.17CE%%rod@example.com,<example@aol.com>
18:34:19 1AiHkY-000Luz-0W 0,3924,rod@example.com,example.demon.co.uk,[192.168.0.1],
         BC3081C2.17D0%%rod@example.com,<paulf@example.co.uk>
18:39:20 1AiHpQ-000MUl-0W 0,3975,rod@example.com,example.demon.co.uk,[192.168.0.1],
         BC3082EF.17D2%%rod@example.com,<hitsch@example.ch>
18:49:48 1AiHzY-0004yv-0Z 0,1986,rod@example.com,example.demon.co.uk,[192.168.0.1],
         BC308560.17D7%%rod@example.com,<jules@example.co.uk>
```

Spotting “spam”

- **Report “too many” failures to deliver**
  - also consider the encoded value of failures
- Ignore “bounces” !
  - have null “< >” return path, these often fail
- Add some further heuristics
  - multiple destinations and all fail
- Unfortunately aol.com never rejects email

Needs (eventually)... 2500 lines of Perl!

```
2004-01-31 08:19:49 242pdr1242pdr1@yahoo.com -> !sksk13@hanmail.net Size=4821
2004-01-31 08:19:55 dixtt6idixtt6i@yahoo.com -> !sksago@hanmail.net Size=4821
2004-01-31 08:20:02 xttiidixttiidi@yahoo.com -> !searchna@hanmail.net Size=4824
2004-01-31 08:27:07 242pdr1242pdr1@yahoo.com -> !sks0607@hanmail.net Size=4823
2004-01-31 08:27:16 2sp4ia22sp4ia2@yahoo.com -> !seamg@hanmail.net Size=4814
2004-01-31 08:27:21 xttiidixttiidi@yahoo.com -> !seain0325@hanmail.net Size=4830
2004-01-31 14:13:13 dixtiridixtiri@yahoo.com -> !skoin@hanmail.net Size=4815
2004-01-31 14:13:18 2pdr1242pdr124@yahoo.com -> !sea4548@hanmail.net Size=4822
2004-01-31 14:13:33 ap2s2iiap2s2ii@yahoo.com -> !skl2035@hanmail.net Size=4823
2004-01-31 14:13:35 242pdr1242pdr1@yahoo.com -> !skko16@hanmail.net Size=4821
2004-01-31 14:13:38 2sp4ia22sp4ia2@yahoo.com -> !se3127@hanmail.net Size=4820
2004-01-31 14:13:55 a22sp4ia22sp4i@yahoo.com -> !skj13@hanmail.net Size=4815
2004-01-31 14:13:59 dixtsaidixtsai@yahoo.com -> !skinart0743@hanmail.net Size=4835
2004-01-31 14:14:09 dixtt6idixtt6i@yahoo.com -> !skid-row0228@hanmail.net Size=4841
2004-01-31 14:14:11 xttiidixttiidi@yahoo.com -> !sds1443@hanmail.net Size=4822
2004-01-31 14:14:20 dixtiridixtiri@yahoo.com -> !skhhks@hanmail.net Size=4821
2004-01-31 14:14:24 dixtsaidixtsai@yahoo.com -> !skh6755@hanmail.net Size=4823
2004-01-31 14:14:31 xttiidixttiidi@yahoo.com -> !sdk5269@hanmail.net Size=4822
2004-01-31 14:14:38 2pa2pis2pa2pis@yahoo.com -> !skftlwkd@hanmail.net Size=4825
2004-01-31 14:14:40 ap2s2iiap2s2ii@yahoo.com -> !skfmfrkwuqhk@hanmail.net Size=4841
2004-01-31 14:14:47 dixttiidixttii@yahoo.com -> !sketco@hanmail.net Size=4821
2004-01-31 14:14:52 xtsaidixtsaidi@yahoo.com -> !sdh1005@hanmail.net Size=4822
2004-01-31 14:14:59 uttiyryuttiyry@yahoo.com -> !sdfam@hanmail.net Size=4814
2004-01-31 14:15:02 a22sp4ia22sp4i@yahoo.com -> !ske08@hanmail.net Size=4815
2004-01-31 14:15:08 2sp4ia22sp4ia2@yahoo.com -> !sdd74@hanmail.net Size=4814
```

Viruses

- Common for mass mailing “worms” to use address book (mainly valid addresses)
- Recent trend towards using contents of email storage, browser cache and (Swen) accessing Usenet servers via NNTP
  - so many addresses now invalid or badly formed
- **So virus infections are also detected**
  - ⚐ but HELOs are a dead give-away!

```
HOST = example.demon.co.uk, HELO = Pozdds
2004-01-23 15:37:09 info@example.demon.co.uk -> ShiongJoo@example.com.my Size=135574

HOST = example.demon.co.uk, HELO = Vgkgqyldf
2004-01-23 15:37:15 info@example.demon.co.uk -> siewmei@example.com.my Size=134898

HOST = example.demon.co.uk, HELO = Hprnljq
2004-01-23 15:37:20 info@example.demon.co.uk -> !Veronique_Prigaux@examplegroup Size=134896

HOST = example.demon.co.uk, HELO = Rbg
2004-01-23 15:37:25 info@example.demon.co.uk -> alanmassow@example.net Size=137550

HOST = example.demon.co.uk, HELO = Enrlnsfdi
2004-01-23 15:37:31 info@example.demon.co.uk -> alison@example.freeserve.co.uk Size=137601

HOST = example.demon.co.uk, HELO = Ukgod
2004-01-23 15:37:38 info@example.demon.co.uk -> ashley@example.co.uk Size=135068

HOST = example.demon.co.uk, HELO = Vsrzo
2004-01-23 15:37:44 info@example.demon.co.uk -> !bestfit@example.net Size=134773

HOST = example.demon.co.uk, HELO = Jmeoqa
2004-01-23 15:37:49 info@example.demon.co.uk -> !cechala@example.net.co Size=134663

HOST = example.demon.co.uk, HELO = Tian
2004-01-23 15:37:54 info@example.demon.co.uk -> Christine@example-recruitment.com Size=133830
```

Validation

- Against remote site reports
  - scheme misses low volume virus infections
  - AOL has their own (high noise) feedback
  - conclusion: we have a very successful detector
- Against traditional “top 50” measures
  - shows we are failing to deal with loops
- BUT excessive “false positives”
  - too expensive to handle ~500 a day

Mail loops

- Simple loop spotted by smarthost
  - **“`Too many "Received" headers`”**
- Simple loop spotted by customer
  - repeated mail ID
  - if no mail ID, then same source & destination
- Responding to Mailer-Daemons
  - may not be a loop at present, but will be!

⚐ **viz: relatively easy to detect mail loops**

```
Message ID = E1Aaxoi-0008C4-I2.2003-12-29-13-52-21@imailg3.svr.pol.co.uk
2004-01-05 09:08:57 dick@gochampion.com -> enquiries@example.freeserve.co.uk Size=4476
2004-01-05 09:22:21 dick@gochampion.com -> enquiries@example.freeserve.co.uk Size=5365
2004-01-05 09:48:32 dick@gochampion.com -> enquiries@example.freeserve.co.uk Size=6254
2004-01-05 10:03:15 dick@gochampion.com -> enquiries@example.freeserve.co.uk Size=7143

Message ID = E1AaM8L-0007ZR-GV.2003-12-27-21-38-05@imailm1.svr.pol.co.uk
2004-01-05 09:08:57 coltons@helicon.net -> enquiries@example.freeserve.co.uk Size=4503
2004-01-05 09:22:22 coltons@helicon.net -> enquiries@example.freeserve.co.uk Size=5392
2004-01-05 09:48:22 coltons@helicon.net -> enquiries@example.freeserve.co.uk Size=6281
2004-01-05 10:03:17 coltons@helicon.net -> enquiries@example.freeserve.co.uk Size=7170

Message ID = 5g-mg$2ex0r7$t25v0w@c5m9o97
2004-01-05 09:08:54 h58wjuo@juno.com -> enquiries@example.freeserve.co.uk Size=4469
2004-01-05 09:22:23 h58wjuo@juno.com -> enquiries@example.freeserve.co.uk Size=5355
2004-01-05 09:48:22 h58wjuo@juno.com -> enquiries@example.freeserve.co.uk Size=6241

Message ID = BAY9-DAV48ylqgrKQ2C0001e111@hotmail.com
2004-01-05 09:09:36 mikepage007@hotmail.com -> kirsty@example.freeserve.co.uk Size=2432
2004-01-05 09:22:23 mikepage007@hotmail.com -> kirsty@example.freeserve.co.uk Size=3319
2004-01-05 09:48:21 mikepage007@hotmail.com -> kirsty@example.freeserve.co.uk Size=4206

Message ID = 0w-$i65-zy-g9@99b.nfgy.m.1db
2004-01-05 09:09:38 np3gvp@wild-flower.net -> enquiries@example.freeserve.co.uk Size=2175
2004-01-05 09:22:34 np3gvp@wild-flower.net -> enquiries@example.freeserve.co.uk Size=3067
2004-01-05 09:48:18 np3gvp@wild-flower.net -> enquiries@example.freeserve.co.uk Size=3959
```

False positives

- Bounces
  - customer rejects email without NULL path
    - bounce@
    - postmaster@
    - root@
- Mailing lists
  - if many failures then looks like spammer!
  - fortunately, most spammers change sender

unwise
Other ~~clueless~~ behaviour

- Saying HELO to destination when bounce
- Invalid MAIL FROM
  - `username@pop3.demon.co.uk`
- Vacation programs
  - small message sent back to every spammer
  - but they fail -- and hence look like spam
- Virus rejections
  - dear 3rd party, I’ve just stopped MyDoom

Some statistics I

- 14 day period (1-14 Nov)
  - 76K customers, 16.8 million emails
- 47 open servers
  - zero “false negatives”
  - 87 “false positives”
    - 2 failing senders indicates “spam” (5 is better?)
    - 498 correctly spotted as “clueless” bouncing or mailing lists with high incidence of failures

Some statistics II

- 34 virus infections (mainly Swen)
  - 22 false positives (5+ HELOs at single site)
- 147 email loops
  - 23 false negatives (>1200 emails but not detected as a loop)
  - Considering email size in preference to message ID for “same again” will improve the false negatives considerably

More to come ???

- Real-time processing of reports
  - probably not worthwhile since rejections are key aspect of heuristics & these can be slow
  - real-time assessing of 5xx messages may assist
- Tar-pits
  - once we have ?100 failures then slow down
  - collateral damage on poorly run mailing lists
  - shutting the stable door (but assists reputation)

About those static addresses...

- Have been assuming that customer records can be collated by IP address
  - ie: static IP address
- Many ADSL systems have long term IP address stability (even if DHCP)
  - ie: static on a day-to-day basis
- For dynamic IP (ie dialup) then collating events with similar times looks promising

What next?

- Spammers will evolve!
  - Looking like bounces will be hard to deal with
  - A valid MAIL FROM will be harder to detect
  - Reducing the volume will be harder to spot
- Viruses will evolve!
  - Changing HELO isn’t doing them much good
  - May begin spot nonsense destinations

Darwin was right!

User Modelling [academics woz here!]

- Stolfo et al (Columbia)
  - collect stats on “normal” behaviour
  - get excited when things change
- but ISP customers are often businesses
  - multiple people
  - random decisions to launch mailshots
- Doesn’t look as if single valued statistics would be sufficient to spot anomalies

Dealing with MyDoom

- Mainly detected as “open server” ☹
  - many customer sites suppressed the HELOs
  - multiple emails meant HELO re-used
- But very effective detection ☺
  - spotting just a dozen emails sent from a dialup
  - figures commercially sensitive, so unpublished
  - one week later, ~80% now cleaned up!

Conclusions

- It is always worthwhile to wonder why your own strange behaviour hasn’t been detected
- Spammers & viruses that hide a pattern at the destination make a pattern at the source
- Some simple heuristics currently spot these patterns : with delivery failures being key
- False positives mainly caused by software & users that are being especially clueless ☹

Extrusion Detection

THE END : Any questions ??

UNIVERSITY OF
CAMBRIDGE
Computer Laboratory

thus™