Advanced Network Security

Richard Clayton

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Check Point Course

7 September 2009

Overview

- Is the infrastructure secure ?
  - attacks on DNS
  - attacks on BGP
- ISP log processing
  - using heuristics to detect email spam

All your mailserver
are belong to us

CAUTION

This talk describes possible attacks on Internet infrastructure, especially DNS & BGP. But, not all of these attacks work everywhere, and people may be reluctant to discuss whether they work or not in their part of the real world.

So don’t assume it’s all entirely true!

However, it isn’t entirely false either!

Any mention I make of particular networks, ISPs or countries is merely to make abstract ideas concrete, not an analysis of actual flaws.

NB: Do not try any of this at home (OR at work)

Threat scenario

- I wish to capture a significant amount of incoming email to a major ISP mail server
  - email may contain passwords etc
  - email can be made to contain passwords etc
  - answering email often “proves” identity
  - obvious opportunity to blackmail the ISP, or just trash their reputation as being secure
- Attack should “scale” to many ISPs
  - 0-day exploit on **`sendmail`** not considered here

Resources

- Back bedroom attackers
  - can now have control of a reasonable size botnet
- Criminal entrepreneurs
  - may own (or 0wn!) a smallish ISP in Ruritania
- Organised crime ??
  - simpler for them just to bribe an employee!
- I am NOT assuming that BGP or DNS are too obscure to be attacked effectively

Underlying strategies

- Cannot just steal packets – people notice
  - cf YouTube outage in February 2008 (Pakistan Telecom)
- Accept email, resend to the correct ISP
  - top 50 senders is a give-away, so use botnet
- Reject email end of data with a 4xx response
  - email generally re-delivered after a delay, so suitable for intermittent attacks
- Tunnel SMTP packets to correct place
  - either a peer of target or customer within target

DNS (I): active attacks

- DNS server asks for data
  - attacker supplies incorrect answer first
    - 16 bit identifier is not long enough!
    - but, modern software randomises request port
- Older software is flawed
  - predictable random numbers!
    - or even accepts non-authorised data!
- No-one monitors for attacks
  - however this scales badly, so of limited interest
  - BUT WAIT!

DNS (II): Kaminsky

- Ask for multiple sub-domains (sub1, sub2 etc.)
  - neat way of ensuring resolver always has to ask
- Attacker tries to get their answer in first
  - BUT of course only poisons some obscure sub-domain
- Kaminsky realised could supply NS data as well
  - “in-bailiwick” data (extra info from authoritative server)
  - relied upon for some purposes! So devastating attack!
- Mitigate (only) with lots of entropy (as before)
  - and what of clever servers behind dumb firewalls?
  - only real fix is DNSSEC

DNS (III): phishing

- “Rock-phish” gang spoofed GoDaddy Aug07
  - probably just wanted some cheap domains
  - BUT control of a registrar account permits changes to name server identities
- Registrars for grown-ups will check validity of changes out-of-band, $10 hosting will not
  - significant number of US banks were vulnerable
- Attack vector might also be malware…

DNS (IV): root of trust

- 13 top level name servers (A-M)
  - maximum that will fit in a DNS response
- Included with BIND (etc) as a text file
  - you have to start bootstrapping somewhere!
- L moved from 198.32.64.12 to 199.7.83.42
  - moved 1 Nov 2007 (warnings sent 24 Oct 2007)
  - AS20144 (ICANN) announced route until 2 May 2008
- BUT other AS's announced route
  - Dec 15 (AS42909), Mar 18 (AS 4555), Apr 1 (AS9584)
  - all serving the right thing (through May, we think!)

Attacks on BGP

- Basic idea: announce a /32 for mailserver
  - BGP prefers a “more specific” announcement
- Traffic then flows to Ruritania
  - email contents are available for inspection
- /32 may not propagate, so /24 may be better
  - leads to complexity if other hosts or services on /24
  - hence tunnelling packets back to ISP may be best (and just sniff them as they pass)
- Sniffing possible anyway at other ISPs
  - difference here is scale and remoteness

More specifics...

- Route should not be accepted
  - mnt-lower prevents creation of new route objects
  - so everyone ought to notice that route isn’t valid
  - complexities with multiple registries
- Route may be spotted by monitoring
  - MyASN @ RIPE, Renesys etc
  - note that bogon filtering hides route from owner! and so Best Practice prevents give-away failures

Unauthorised announcements

- Existing route: hope to be a shorter AS path
  - BGP counts AS's to determine preference
  - so more effective in Ruritania than London
- May help to forge origin for peer to accept the route (entirely dependent on filters)
- Once again, monitoring detects wickedness
  - but registry data error-prone and incomplete so can perhaps only consider changes?
  - and of course you need to know all about multi-homed customers! Is this possible?

More BGP Stuff

- RIPE

  MyASN & lots of other initiatives

- Experimental alerting systems

  `http://iar.cs.unm.edu/alerts.php`

  `http://phas.netsec.colostate.edu`

- Anirudh Ramachandran and Nick Feamster

  SIGCOMM 2006: Understanding the Network-Level Behavior of Spammers

SMTP Defence I: encryption

- Opportunistic encryption (RFC3207)
  - uses STARTTLS capability & command
  - negotiate mutually acceptable algorithm
- Plus points:
  - works out of the box for major MTAs
  - only end-points can decrypt the traffic
- Minus points:
  - increases processing load (may not matter)
  - no “man-in-the-middle” protection

SMTP Defence II: authentication

- Check certificates before sending email
  - prevents man-in-the-middle
- Plus points:
  - works out of the box for major MTAs
- Minus points:
  - increases processing load (albeit may not matter)
  - needs a Public Key Infrastructure (or a lot of bilateral arrangements)

Network level defences

- Anti-spoofing filters on customer links
  - motherhood! (but tedious for custom customers)
- Much harder to do on border routers
  - unicast reverse path forwarding (RPF) can help
  - but at IXPs this may not be practicable
- Can check if traffic coming from correct peer
  - straightforward(ish) sFlow/Netflow analysis

Secure DNS/BGP

- Secure DNS almost here
  - some TLDs already signed, more to come
  - unlikely that will be fully deployed for years
  - BUT Kaminsky exploit has given it a huge boost
- Secure BGP(s) experimental at present
  - concerns about performance (cf MD5)
  - concerns about key distribution
  - when will it be stable and inter-working?

Blended attacks

- Some key distribution schemes use DNS
- Attack the DNS and you may be able to compromise systems that are “secure”
- Best use of a BGP attack may be to capture the DNS servers (think long TTL), and then you can go after the mail servers at leisure!
- ...and of course you may just want to DoS
  - so you don’t mind if your attack is noticed

But why not just attack the customer directly?

Customer equipment

- Windows machines may keep name server identities in registry – easy for malware to change
- But in practice, usually set by DHCP
- Hence only need to compromise home routers
  - may have no password at all (and insecure wireless)
  - may be configurable from “the outside”
  - may be insecure, with buffer overflows &c
  - may still have the standard password
- With wireless as well, some researchers postulate an out-of-band worm!

Negligence

- The failure to use reasonable care
- Current test for “duty of care”:
  - harm must be (1) reasonably foreseeable
(2) there must be a relationship of proximity between the plaintiff and defendant and
(3) it must be “fair, just and reasonable” to impose liability
- If one of my attacks is effective on a mailserver, because of firewall failings, are you negligent?
- Short term specific: if your router/firewall makes DNS IP-IDs predictable, are you negligent?

Looking for spam in ISP logs

Email “spam” : key insight

- Lots of spam is to ancient email addresses
- Lots of spam is to invented addresses
- Lots of spam is blocked by remote filters

- Can process server logs to pick out this information. Spam has many delivery failures whereas legitimate email mainly works

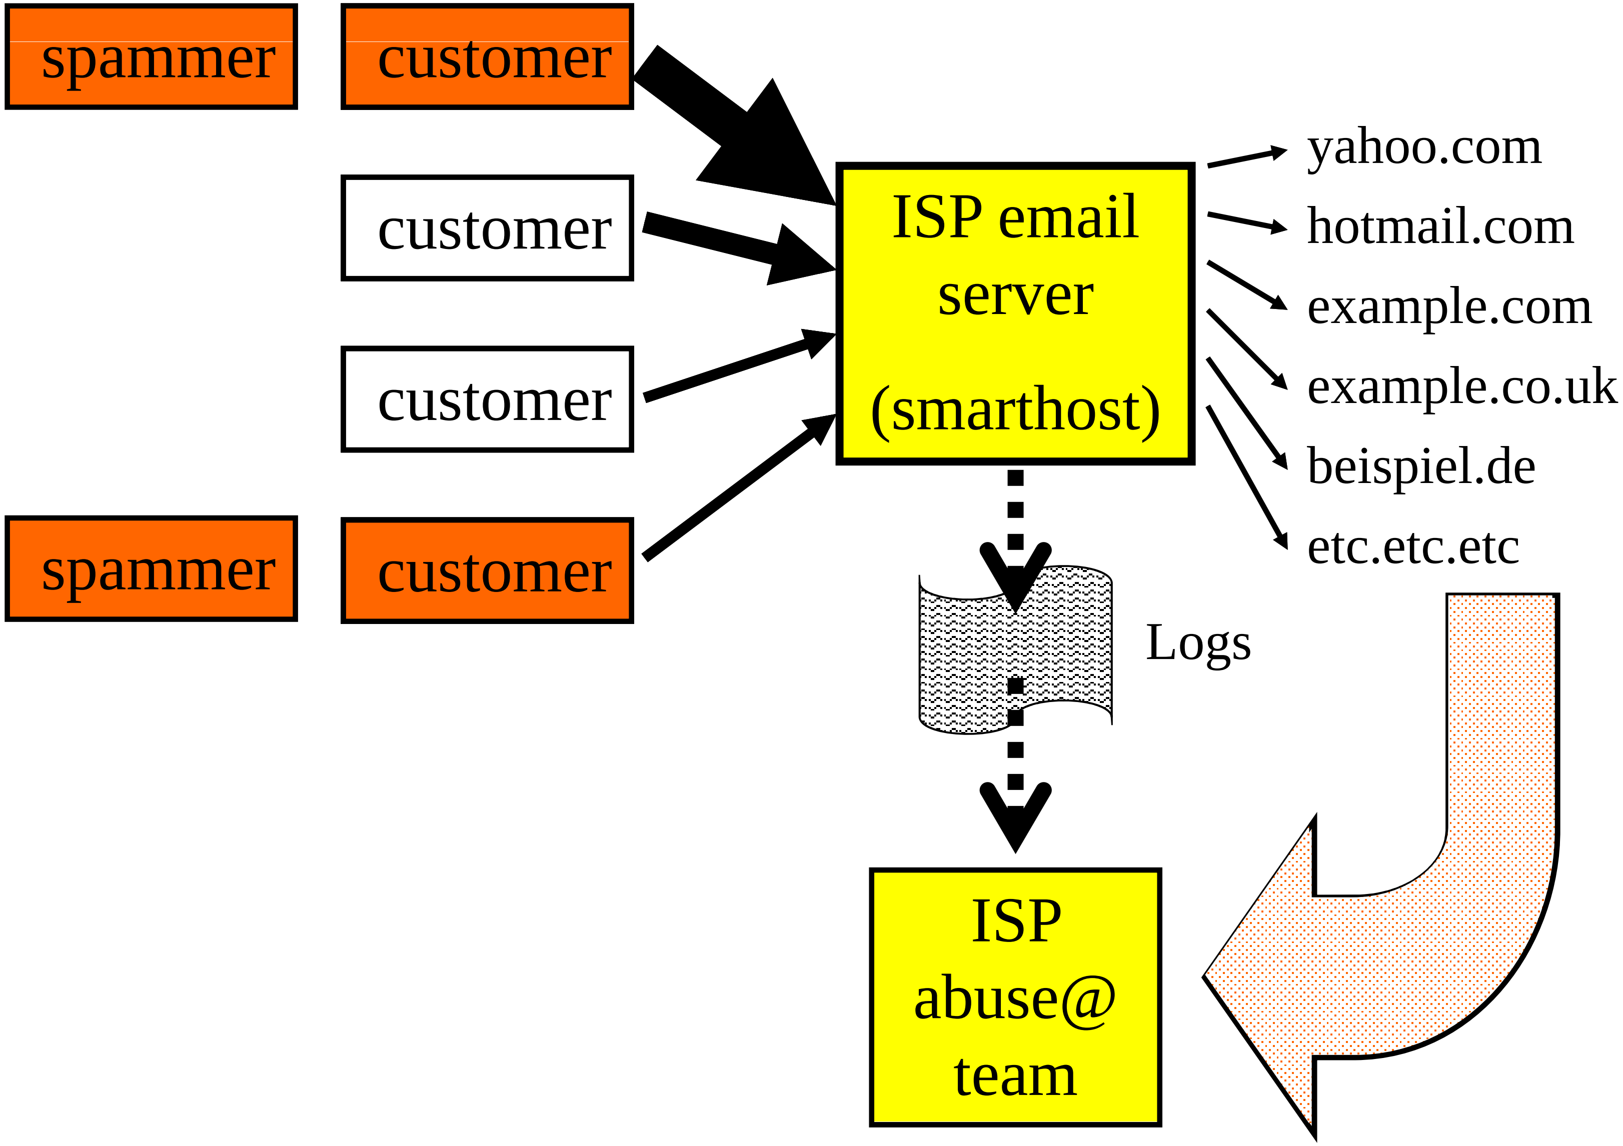
spammer
customer
customer
customer
spammer
customer
ISP email server
(smarthost)
yahoo.com
hotmail.com
example.com
example.co.uk
beispiel.de
etc.etc.etc
Logs
ISP abuse@ team

Log processing heuristics

- **Report “too many” failures to deliver**
  - more than 20 works pretty well
- Ignore “bounces” !
  - have null “< >” return path, these often fail
  - detect rejection daemons without < > paths
- Ignore “mailing lists” (fixed sender)
  - most destinations work, only some fail (10%)
  - more than one “mailing list” is a spam indicator!
- Ignore “forwarding” (fixed destination)
  - multiple forwarding destinations is common

Bonus! also detects viruses

- Common for mass mailing “worms” to use address book (mainly valid addresses)
  - though worms are currently rather out of fashion
- Often remote sites will reject malware

AND, VERY USEFUL!

- Virus authors don’t know how to say HELO
- **So virus infections are also detected**
  - out of fashion, but many still getting infected

```
2007-05-19 10:47:15 vzjwcqk0n@msa.hinet.net          Size=2199
                    !!! 0930456496@yahoo.com
                    !!! 09365874588@fdf.sdfads
                    !!! 0939155631@yahoo.com.yw
                     -> 0931244221@fetnet.net
                     -> 0932132625@pchome.com.tw
2007-05-19 10:50:22 985eubg@msa.hinet.net            Size=2206
                    !!! cy-i88222@ms.cy.edw.tw
                    !!! cynthia0421@1111.com.tw
                     -> cy.tung@msa.hinet.net
                     -> cy3219@hotmail.com
                     -> cy_chiang@hotmail.com
                     -> cyc.aa508@msa.hinet.net
                  and 31 more valid destinations
2007-05-19 10:59:15 4uzdcr@msa.hinet.net             Size=2228
                    !!! peter@syzygia.com.tw
                     -> peter.y@seed.net.tw
                     -> peter.zr.kuo@foxconn.com
                     -> peter548@ms37.hinet.net
                     -> peter62514@yahoo.com.tw
                     -> peter740916@yahoo.com.tw
                     and 44 more valid destinations
```

```
HELO = lrhnow.usa.net

2007-05-19 23:11:22 kwntefsqhi@usa.net                Size= 8339
                     -> ken@example1.demon.co.uk

HELO = lkrw.hotmail.com

2007-05-19 23:11:24 zmjkuzzs@hotmail.com              Size=11340
                     -> ken@example2.demon.co.uk

HELO = pshw.netscape.net

2007-05-19 23:14:52 dscceljzmy@netscape.net           Size= 6122
                     -> steve.xf@example3.demon.co.uk

HELO = zmgp.cs.com

2007-05-19 23:18:06 wmqjympdr@cs.com                  Size= 6925
                     -> kroll@example4.demon.co.uk
```

ISP email handling

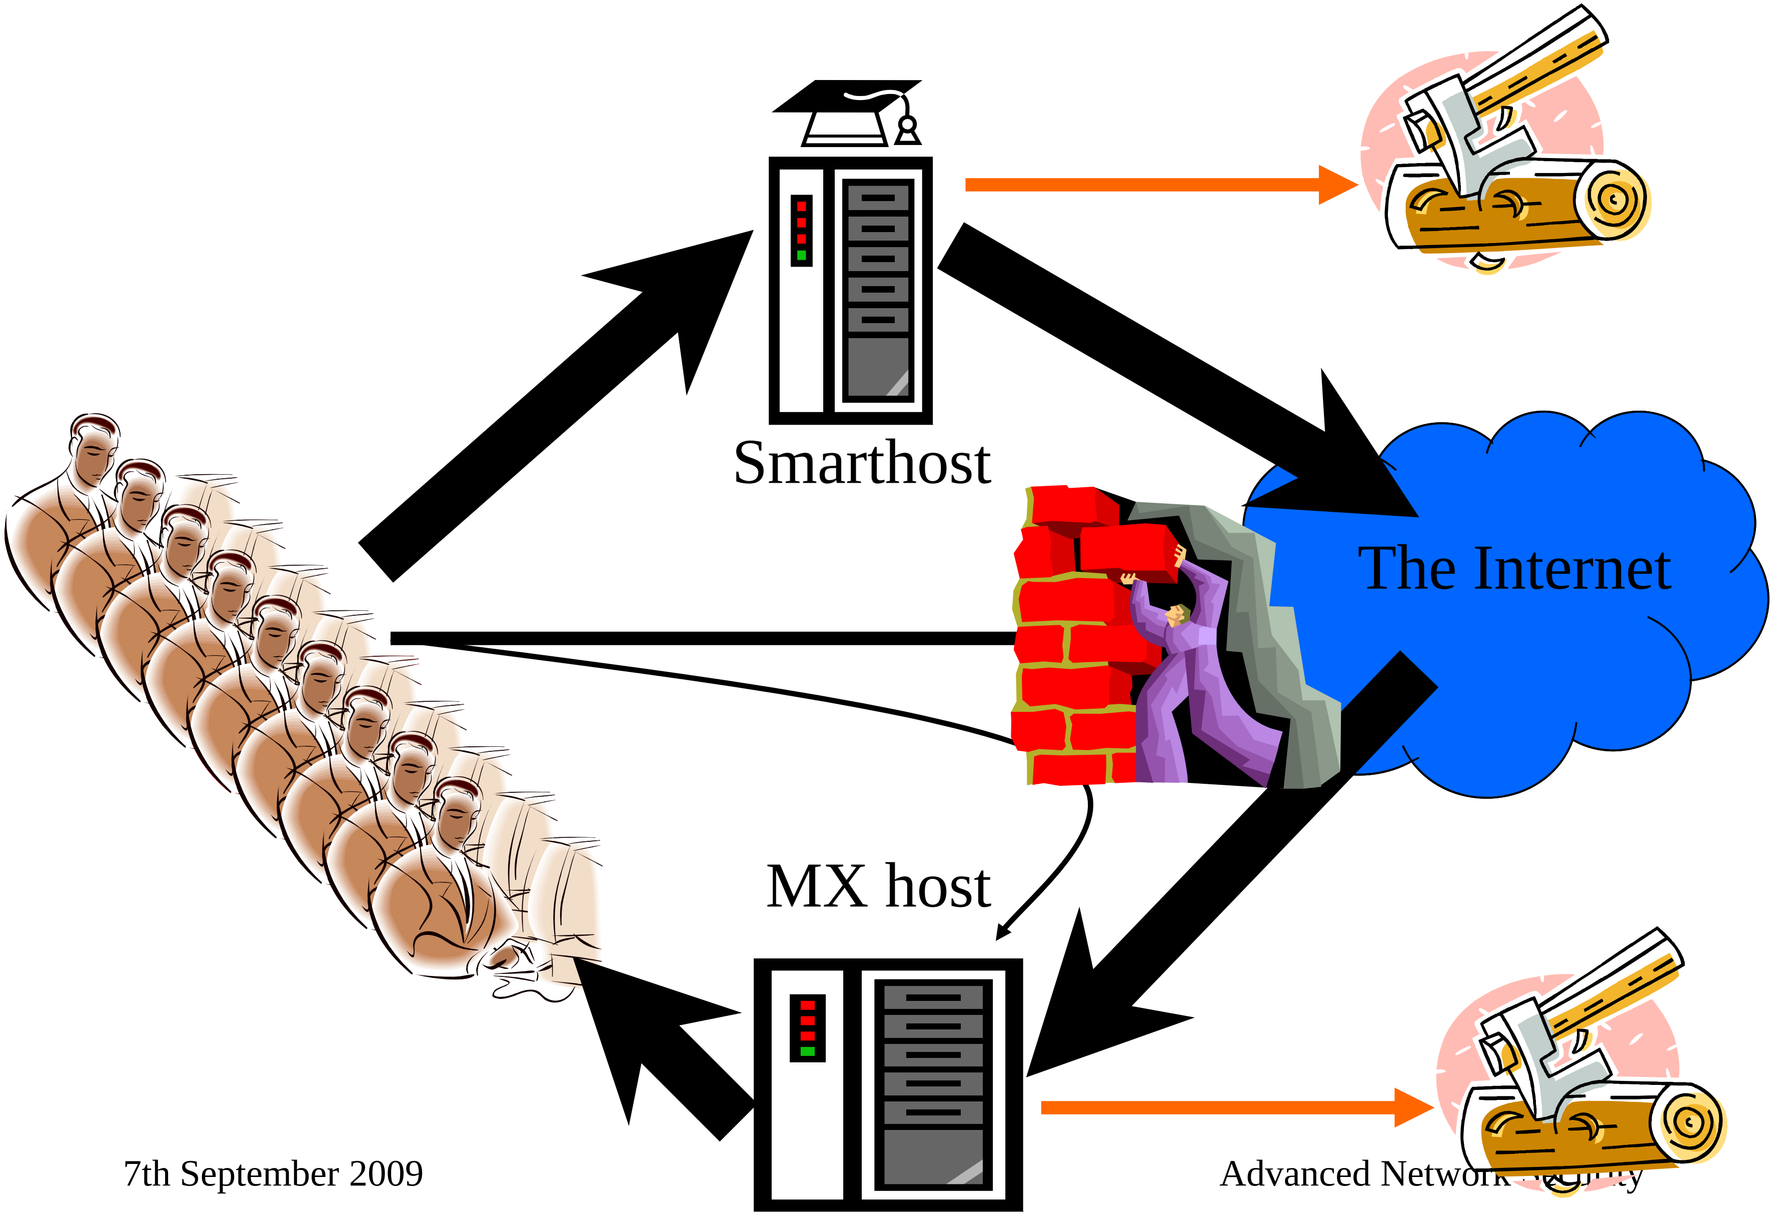

Incoming email

- Some spam runs will also target other customers
  - complex for spammers to avoid this
- Some spammers try and use the smarthost, but using the MX record doesn’t work too well
  - major ISPs don’t do “in” and “out” on the same machine
- Hence processing incoming server logs can locate the spammers who don’t use the smarthost
  - heuristics can in fact be set much more sensitively
  - once again, good at spotting virus activity

Email log processing @ demon

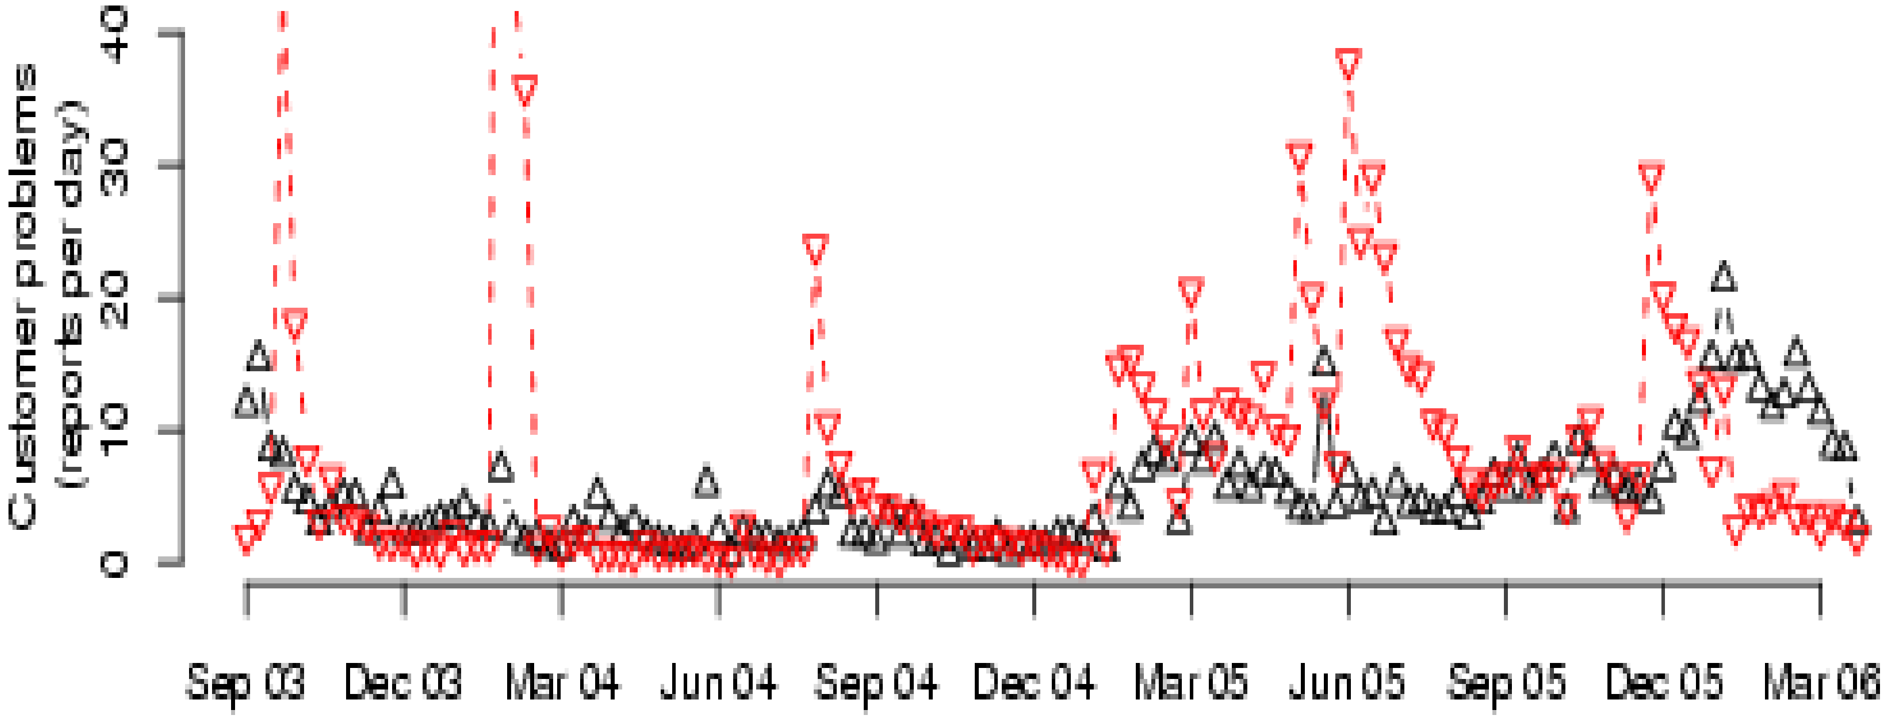

Detection of spam (black) and viruses (red)

Incoming reports (all sources)

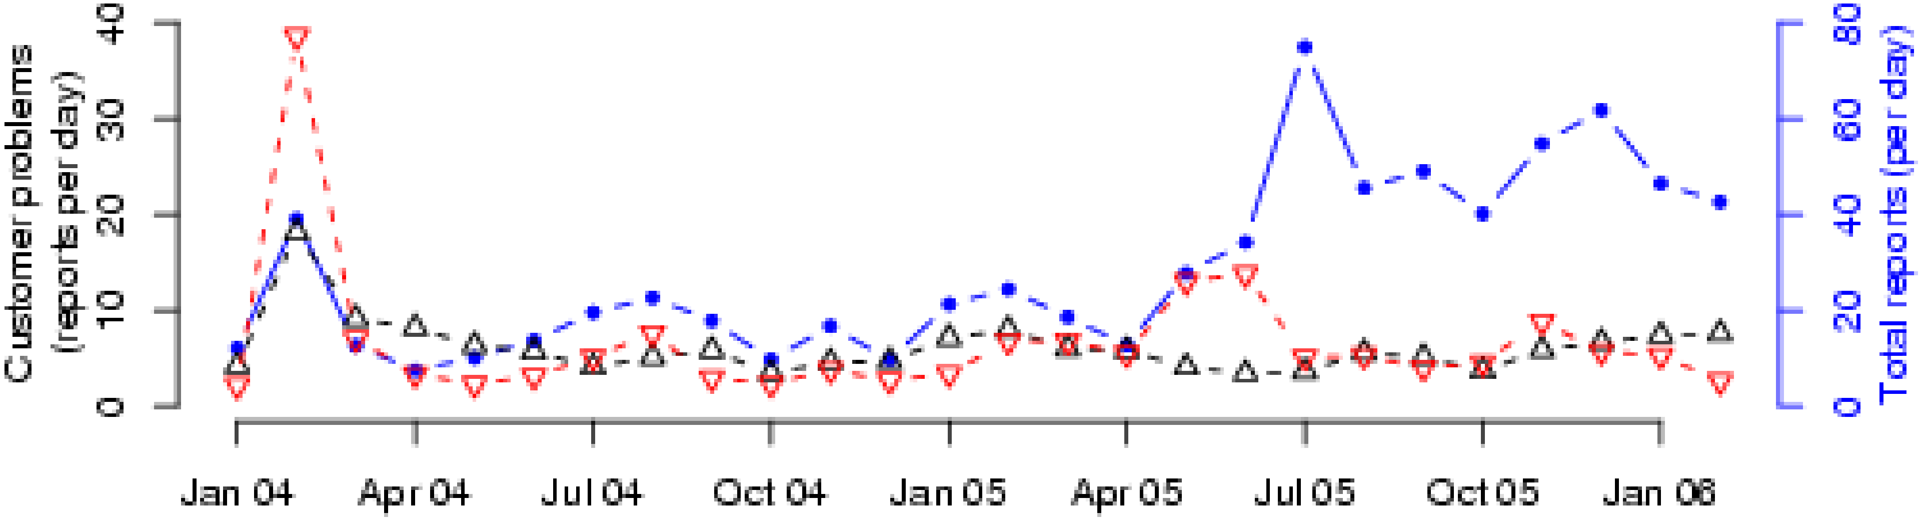

spam (black), viruses (red), reports (blue)

Traffic analysis

- This is a specific example of a general technique called “traffic analysis” which permits analysis of activity without access to the content
- The spooks have done it for ages, but is now getting significant traction in open community

- This leads to an even more general principle:

  “It’s hard to make one thing look like another”

  especially when attacker doesn’t know exactly (for your chosen measurement) what “normal” looks like

Advanced Network Security

Richard Clayton

http://www.lightbluetouchpaper.org

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Check Point Course

7 September 2009