

Spotting spam in sampled sFlow

Richard Clayton

WACI, Cambridge MA, 3rd October 2007




**UNIVERSITY OF
CAMBRIDGE**
Computer Laboratory



Demon

ISP email handling



Heuristics for log processing


- Simple heuristics work really well
- Key measure is failures to deliver
 - addresses are old/constructed/blocked
- Multiple HELO lines very common in spam
- Look for outgoing email to the Internet
- Pay attention to spam filter results
 - but need to discount forwarding

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

Incoming reports (all sources)



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

spamHINTS research project



LINX samples 1 in 2000 packets
(using sFlow) and makes the port 25
traffic available for analysis...


Basic idea

- Spam doesn't look like normal email, so it can be detected by analysing the traffic patterns
- Big benefits if this can be shown to work, only evasion technique is to look more like normal email (and send less traffic)
- Running this at a major IXP (LINX) improves accuracy & permits amortisation of costs (and development) across the whole industry
- Port 25 is an OK discriminator !


Challenges

- No content data
 - part of agreement, so had to be removed
- sFlow is sampled
 - sampling is of packets, my data is then filtered from that (but large numbers should avoid bias)
- Only Foundry ring currently instrumented
 - Extreme implementation not ready for prime time
- Some private peering (so flows missed)

Known “open server”




Another known “open server”




Look for excessive variation

- Look at number of hours active compared with number of four hour blocks active
- Use incoming email to Demon to pick out senders of spam and hence annotate them as good or bad...
- ... did analysis for a large ISP, but problem is that “if it sends, it’s bad”. Nevertheless...



Spamminess vs hours of activity for IPs active in 5 of 6 possible 4 hour periods



So work continues...

- sFlow data will always be useful to feed back ongoing activity to abuse teams
- Analysis may improve when both rings instrumented and when data available in real-time (so can compare historic data)
- Still to consider variations (and lack of variations) in destination as well as time

A related approach

Filtering Spam with Behavioral Blacklisting

Anirudh Ramachandran, Nick Feamster, and Santosh Vempala
to appear at upcoming: ACM CCS (Oct 29 – Nov 2 2007)

Uses a spectral clustering algorithm to try and divide sending IPs into groups. Assesses sending (per IP) to email addresses within 150 domain names (viz: SMTP log level data). Idea is that spammer will target same sets of domains, but from a new IP address.

Summary

- Attempting traffic analysis on sampled sFlow
- Sampling means data rates are rather low
- Labelling of IP addresses also tricky
- Much more work needed on good distinguishers
- But would be really useful if it worked ☺☺☺

<http://www.cl.cam.ac.uk/~rnc1>

CEAS papers: <http://www.ceas.cc>

2004: Stopping spam by extrusion detection

2005: Examining incoming server logs

2006: Early results from spamHINTS

2007: Email traffic: A qualitative snapshot



**UNIVERSITY OF
CAMBRIDGE**
Computer Laboratory



Demon