Using Early Results from the 'spamHINTS' Project to Estimate an ISP Abuse Team's Task

Richard Clayton
Computer Laboratory, University of Cambridge
JJ Thomson Avenue, CAMBRIDGE, CB3 0FD, United Kingdom
richard.clayton@cl.cam.ac.uk

ABSTRACT

ISPs operate "abuse" teams to deal with reports of inappropriate email being sent by their customers. Currently, the majority of this work is dealing with insecure systems that have become infected with viruses or that have been hijacked by the senders of email "spam". This paper examines the performance of an abuse team at a large UK ISP over the past few years, and shows that email log processing tools have provided significant improvements in their efficiency. A new email measurement system called spamHINTS, using sampled sFlow packet header data from a major Internet exchange, is currently under development. Early results from this monitoring suggest that the ISP abuse team needs to step up their activity by an order of magnitude to get on top of their problem.

1. INTRODUCTION

The sending of bulk unsolicited email ("spam") has been a major problem for over a decade. Originally, spammers would sign up with an Internet Service Provider (ISP) and send as much email as possible, until their activity was detected and their account was closed. Spammers then sent their email via "open relays" which sent email without checking for authorisation, so the ISPs had to educate customers about mail server configuration. More recently, spam has been sent via "trojan" software on end-user machines; and once again the ISPs have found themselves educating customers – on the value of anti-virus software, applying software patches, and other security measures.

Demon Internet is a large (200 000 customer) ISP in the United Kingdom. Its customer base is a mixture of consumers, small businesses and some large corporations. Every month, its abuse team handles many thousands of reports of spam and virus traffic sent by customers (almost invariably through insecurity rather than intention). Historically, these reports came from individuals who had received unwelcome email, but other sources now provide the vast majority.

Demon Internet also operates a number of proactive measures to detect customer problems. For several years end-user machines have been regularly "scanned" to determine if they are "open relays", which is a relatively common practice. More unusually, since 2003 Demon Internet has been analysing outgoing email server (smarthost) logs to detect the traffic patterns suggestive of the sending of spam [1]. Also, since early 2005, incoming email server logs have been examined to spot customers who send spam "directly" to other customers, or even themselves [2]. A review of the success of these schemes is presented in Section 2 below.

After many years of effort in dealing with reports, it might be expected that Demon Internet would have a relatively "secure" set of customers. Open relays have been reconfigured, anti-virus software has been installed on systems that were once infected, and trojans have been detected and removed. The general level of complaints from third parties and the "blacklisting" of the mail servers have reduced considerably, so there is a feeling that the problem, if not solved, is at least under control.

In the past year or so, a number of automated systems at other ISPs have started to operate (at AOL, Earthlink, Road Runner, etc). It is increasingly a precondition for the continued acceptance of other, legitimate, email that reports of spam are accepted and dealt with. The statistics from these systems are presented in Section 2.1 below, and these tend to reinforce the view that there are relatively small numbers of Demon Internet customers with problems.

As previously reported [2], the number of customers with a spam/virus problem can be estimated by examining the incoming email traffic from the rest of Internet – and assuming that the local customer base is no different. An up-to-date version of this calculation is presented in Section 3 below.

However, a new email activity detection system, currently under development, permits far more accurate detection of ISP customers with spam and virus sending problems. It processes traffic data samples from the LINX (one of the largest Internet Exchange points in the world). Early results from this system are presented in Section 4. Despite the cosy picture created by other estimates, they appear to indicate that only a small proportion of Demon Internet's problem customers are currently being reported to the abuse team.

2. NUMBER OF CUSTOMER PROBLEMS

For most of 2004 the main source of reports to the Demon Internet abuse team was the internally operated mail server log processing system. Figure 1 shows the average daily number of reports of both spam ($\triangle$) and virus ($\triangledown$) sources generated on a week by week basis starting in September 2003 – when the outgoing log processing was introduced – up to the present. As can be seen, after catching up with an initial backlog of customer problems, the detection levels settled down to a handful of reports a day.

From early 2005 onwards, incoming server logs were analysed, raising the reporting levels to around 8 customers per

CEAS 2006 - Third Conference on Email and Anti-Spam, July 27-28, 2006, Mountain View, California USA

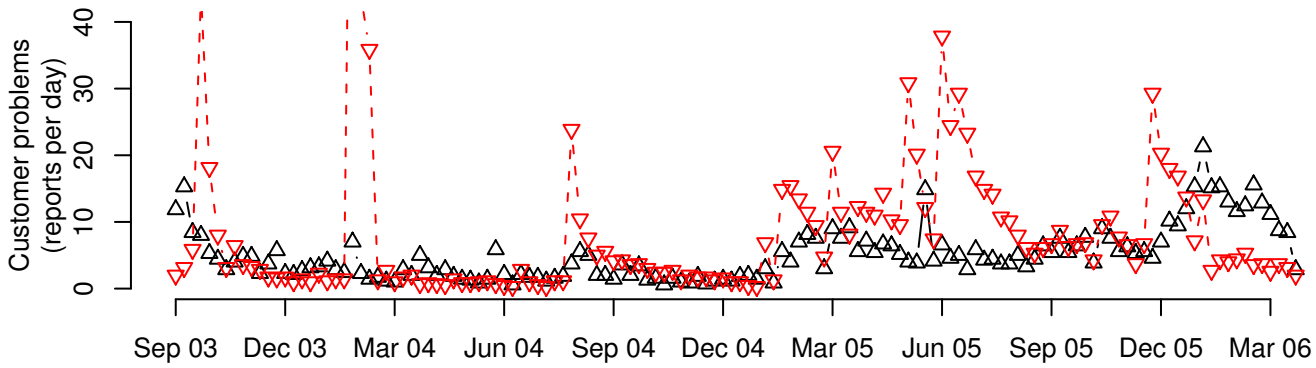

Figure 1: Demon Internet customers detected by log processing as sending spam (△) and viruses(▽)

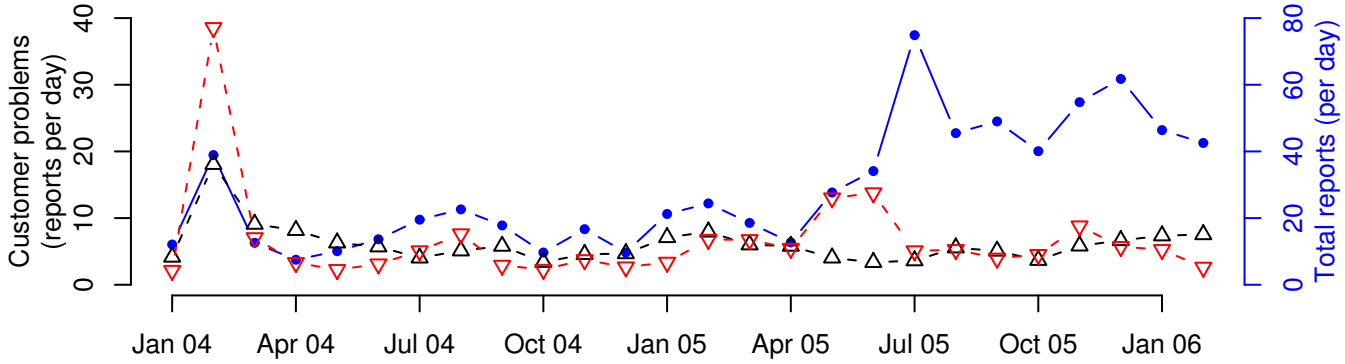

Figure 2: Overall totals of reports (•); and the spam senders (△) and virus senders (▽) detected thereby

day. The big increase (to almost 20 a day) in early 2006 reflects a reworking of the heuristics, in particular, picking out systems that issue an SMTP protocol [5] HELO command with the receiving machine's name or IP address (the correct command would use the sender's identity).

New behaviour by some spam sending programs has also contributed to the log processing's effectiveness. Instead of sending email directly, expecting it to be blocked, they try to use the ISP's smarthost. However, they are unaware of the identity of this machine – and assume that outgoing email can be sent via the machine used for incoming email, which they can easily locate by means of an "MX lookup" for the machine they have hijacked. Since Demon Internet's incoming mail system machines do not accept outgoing email, this "send-to-MX" behaviour is easy to detect.

The graph in Figure 1 also shows that virus reports are generally of the same level as the spam reports – except when new viruses come along. The peak in September 2003 was "Swen", in late January 2004 "MyDoom" (one week, averaging over 100 a day), July–August 2004 the large number of "MyDoom/NetSky" variants and so on.

The overall picture is of an efficient abuse team, initially receiving a large number of reports as a new detection technology (or a new virus) is introduced, who deal with the customers with a problem – and fairly soon detection returns to a steady state.

2.1 Reports from Other ISPs

The total reports to the abuse team is illustrated in Figure 2. These figures are collated monthly, so the peaks of email virus activity (which only last a few days) are less pronounced. The upper line on the graph (joining the • points) is the total number of reports received from all sources (and is plotted on the right-hand axes, viz: at half the scale).

As can be seen, in 2004 and the first half of 2005, most of the reports resulted in the identification of a customer with a problem. Since these figures also include the internally generated reports, it is apparent that that reports from remote sites are only a small part of the abuse team's workload.

However, there is a significant change in mid-2005, when the totals increase considerably. This was when reports started to be accepted in bulk from other ISPs, so as to ensure ongoing connectivity for Demon Internet customers. We note that the detection rate is hardly affected by the very significant number of reports, indicating that there is merely over-reporting of the same problems.

3. ESTIMATING THE SPAM PROBLEM

In May 2005 the US Federal Trade Commission (FTC) wrote to 3 000 ISPs worldwide [4] asking for outgoing customer traffic on port 25 (the TCP port used for SMTP) to be blocked, so that insecure and infected end-user machines would be unable to send email, except via the ISP smarthost. Although this blocking has been widely implemented by consumer ISPs in North America, it is relatively rare elsewhere. In the UK, most ISPs have a mixture of consumer and business customers and while blocking port 25 for consumers might be acceptable, blocking the direct sending of email by businesses is unattractive. Since many ISPs are unsure of the exact status of every customer, it is simplest to

Table 1: Top 20 Spam Sources (by AS)

Count	AS	Description and Country	Ratio
80 319	AS4134	CHINANET (CN)	1.8
75 980	AS4766	Korea Telecom (KR)	1.6
47 578	AS4812	China Telecom (CN)	1.7
18 683	AS9318	Hanaro Telecom (KR)	1.9
12 609	AS4837	CNC (CN)	2.6
5 792	AS12322	Proxad (FR)	2.6
4 941	AS3786	Dacom Corporation (KR)	1.5
4 779	AS7738	TeleBahia (BR)	4.7
3 929	AS9277	Thrunet (KR)	1.6
3 911	AS3320	Deutsche Telecom (DE)	6.4
3 910	AS3462	Hinet (TW)	4.0
3 814	AS3215	France Telecom (FR)	7.4
3 286	AS4788	TM Net (MY)	2.0
3 250	AS4814	CNCGroup (CN)	2.2
3 074	AS19262	Verizon (US)	4.5
2 898	AS4670	Shinbiro (KR)	1.7
2 837	AS8167	Telesc (BR)	5.4
2 532	AS6327	Shaw Communications (CA)	2.5
2 443	AS16338	Auna Telecom (SP)	3.2
2 435	AS3269	Telecom Italia (IT)	5.0

Table 2: Top 20 Virus Sources (by AS)

Count	AS	Description and Country	Ratio
3 172	AS2856	BTnet (UK)	5.4
2 130	AS4134	CHINANET (CN)	1.8
1 727	AS9105	Tiscali UK (UK)	4.7
1 453	AS3320	Deutsche Telekom (DE)	6.4
1 420	AS5089	NTL (UK)	4.6
1 298	AS4766	Korea Telecom (KR)	1.6
936	AS9121	TTNet (TU)	5.2
846	AS9318	Hanaro Telecom (KR)	1.9
782	AS4837	CNC (CN)	2.6
780	AS3215	France Telecom (FR)	7.4
775	AS3352	Telefonica Data Espana (ES)	4.5
757	AS3269	Telecom Italia (IT)	5.0
602	AS5462	Telewest (UK)	3.8
524	AS3462	Hinet (TW)	4.0
495	AS9498	Bharti BT Internet (IN)	5.1
444	AS4755	Videsh Sanchar Nigam Ltd (IN)	5.6
383	AS7132	SBC Internet Service (US)	8.0
346	AS4788	TM Net (MY)	2.0
342	AS5617	TPNet (PO)	7.1
227	AS1267	Infostrada (IT)	12

avoid any blocking; so apart from a few consumer-oriented ISPs the FTC request has mainly been ignored.

If outgoing email is not blocked, a highly relevant question is how much damage is this actually causing? That is, how much spam (and virus) activity is being missed? Any email that goes via the outgoing smarthost is being logged, and the lack of external reports suggests that the log processing heuristics are adequate. When email is sent "direct" to its destination, avoiding the smarthost, incoming server log processing will detect a proportion of the problem because some customers will send email to other customers and hence be spotted. An estimate of this proportion can be calculated by examining what the detection system thinks of the email that is arriving from other ISPs.

3.1 Current Activity

The current version of the log processing system was run on the incoming email traffic (55 900 996 emails) for the first nine days (1^{st}–9^{th}) of May 2006. The system detected 25 997 845 spam emails from 446 157 sources and 1 213 715 virus emails from 45 632 sources. The top 20 ASs (Autonomous Systems – effectively ISPs) from which the incoming spam arrived are listed in Table 1. The numbers detected are somewhat higher than the equivalent data reported in [2] in 2005, but this reflects better heuristics far more than increased activity.

Similarly, the main AS sources of email virus traffic (once again, better detected than in [2]) are shown in Table 2. It can be noted that there are rather more European sources in the virus top 20 than the spam top 20 – reflecting the spreading mechanisms employed.

3.2 Current Estimates

Following the methodology of [2], we consider these major sources of problem email and determined what proportion of the senders within each AS were detected, using the log processing heuristics, to have some sort of abuse problem. The headline "ratio" figures are given in tables 1 and 2. As can be seen, generally between a quarter (ratio=4) and a half (ratio=2) of all the IP addresses that were sending email to Demon Internet were being detected as the senders of spam (or, more rarely, viruses). Since almost all of this email will not have been legitimate, this shows a very significant improvement in the detection technology over the 2005 data, when the "detection ratio" was in the range 30–954.

This suggests that the current detection level of about 20 Demon Internet customer problems per day reflects an underlying population of 40–80 customers who actually have a problem, which is of course quite reassuring.

However, the number of emails per remote site is low, and many of the heuristics will not trigger if the number of emails is less than four or five – so it might be more realistic to assume that *all* incoming email was bad! Typically, between 0.25% to 1.25% of IP address space allocated to these ASs is originating email. Applying these percentages to Demon Internet indicates that 500 to 2 500 senders could have a problem – with the hope being that Demon Internet's long history in dealing with abuse problems will mean that the value is towards the lower end of this range. However, this may not be the case, as we shall now see.

4. THE SPAMHINTS PROJECT

The London Internet Exchange (LINX) is an extremely busy Internet Exchange Point (IXP). It operates two disjoint "rings" running over multiple 10 Gbit Ethernet links between a number of buildings in London's Docklands area. More than 220 member ISPs connect to these rings and exchange Internet traffic over the LINX infrastructure. LINX also provides "private peering" infrastructure for direct ISP to ISP connections. The total traffic over the two rings and the private peering links currently averages just over 100 Gbit/s and peaks at 130 Gbit/s.

In order to provide traffic statistics to their members, LINX is in the process of rolling out a sampled sFlow traffic data scheme. Eventually,[1] one in every 2048 packets that

[1]At this early stage, only one ring produces sFlow data.

enters the two public rings will be sampled and an sFlow [6] record will be made of the IP/TCP/UDP header information. This permits analysis of the source and destination of the packet and what protocol was being used – but no "content" is made available.

LINX have agreed to permit analysis of the email protocol packets (on port `tcp/25`) from the sampled sFlow data. The analysis is being done by the spamHINTS project – HINTS stands for "Happily It's Not The Same" and refers to the analysis of traffic data to pick out the patterns that are indicative of spam. In the longer term it is intended to provide this type of analysis to LINX members in real time, but even at a very early stage it is possible to use the data to determine which machines are sending and receiving email.

On Wed, 10th May 2006 (immediately after the previous data was collected) the sampled sFlow data yielded 9 294 018 SMTP packets (with source or destination port `tcp/25`). Analysing this 24 hours of data showed that email was being sent by 524 244 client machines towards 357 024 servers (or just 492 928 clients and 306 469 servers if SYN and RST packets are ignored). The top 20 ASs sorted by the number of IP addresses sending email are shown in Table 3 (LINX confidentiality conventions prevent a precise identification of the ASs). The disparities with the other data are the subject of further research, but the most likely explanation is the lack of data from the second LINX ring.

Table 3: Top 20 Email Sources (by AS)

IP addresses	Country of origin
20 255	United Kingdom
18 687	Turkey
16 213	Korea
11 349	Spain
10 861	Italy
9 708	Korea
7 596	France
6 834	AS2529, Demon Internet
6 211	Portugal
6 183	United Kingdom
5 334	India
5 026	Germany
4 831	Middle East
4 765	United Kingdom
4 400	United States
4 226	India
4 217	China
4 168	United Kingdom
4 125	France
3 950	United Kingdom

As can be seen, there are 6 834 clients within the Demon Internet AS. Because the data is sampled (one packet in 2048) this is an underestimate – although it is very likely that all the most active systems (handling several hundred emails a day) will have been detected.

Any machine that both sent and received emails (acting as both client and server) was probably intentionally running email software. If we exclude these machines, this leaves 3 485 customers who just sent email – and experience suggests that in many cases they will be unaware they are doing so, viz: they are sending spam or virus infected email. This is a rather higher figure than any of the estimates above. An obvious explanation for the disparity with the incoming server log processing figure (which yielded an estimate of 40–80 customers) is that spam is not sent entirely randomly, but a conscious effort is being made by the spammers to select the destinations the spam is sent to – avoiding "nearby" machines – and hence the assumptions underlying the earlier estimate are not valid.

There is one further source of useful information in assessing whether email clients might be legitimate. Demon Internet maintains a list of IP addresses for customers who wish to be able to send email directly to AOL. This is because AOL routinely blocks incoming email from ISP customer address space unless special arrangements have been made. However, only 143 of the 3 485 sender addresses appear on this "whitelist" so it makes no substantive difference to the conclusion (in fact it serves to reinforce it) that over 3 000 Demon Internet customers are currently likely to have an email sending problem.

5. CONCLUSIONS

Examining long-term statistics for the activity of the Demon Internet abuse team demonstrates that the email log processing systems are providing most of their email-related workload. The bulk sending of reports from other ISPs has significantly increased the number of reports to be processed, without making very much impact on the number of customer problems being detected.

Using the log processing tools to analyse the activity of customers of other ISPs indicates that these tools are at least 10 times better than previously reported at detecting problems. This analysis leads to an estimate that at any given time there are about 80 Demon Internet customers with a spam/virus problem and about a quarter of these are being detected and dealt with.

However, examining packet level traffic data from the spamHINTS project at the LINX is a potentially a far more accurate way of assessing the problem, and this paints a far gloomier picture. There may be more than 3 000 customers who are sending email unknowingly – tackling this within a month would involve a workload of 100 new cases a day, an order of magnitude more work than the abuse team currently have to tackle.

The spamHINTS project is still at a very early stage and only provides simple lists of email sources and destinations. It is being developed to pick out the patterns in time, destination and size that will be indicative of the sending of spam or viruses, rather than the sending of legitimate email. Once that processing is in place it will be possible to refine the estimates of compromised customers – and, more importantly, start to generate useful reports to ISP abuse teams not only at Demon Internet, but elsewhere in the world as well. It looks as if they may have a fair amount to do.

6. ACKNOWLEDGMENTS

The spamHINTS project is financially supported by Intel Research and endorsed by the membership of LINX. We wish to also recognise the vital support of Demon Internet in developing systems for accessing the information buried in email server logs. The patience and professionalism of their abuse team in dealing with the automatically generated reports has also made a significant contribution to this work.

7. REFERENCES

[1] R. Clayton. Stopping spam by extrusion detection, First Conference on Email and Anti-Spam (CEAS 2004), July 2004.

[2] R. Clayton. Stopping outgoing spam by examining incoming server logs, Second Conference on Email and Anti-Spam (CEAS 2005), July 2005.

[3] D. Crocker: Mailbox names for common services, roles and functions. RFC2142, IETF, May 1997.

[4] Federal Trade Commission: FTC, partners launch campaign against spam "Zombies". Press release, FTC, 24 May 2005.

[5] J. Klensin: Simple mail transfer protocol. RFC2821, IETF, April 2001.

[6] P. Phaal and M. Levine: sFlow version 5. sFlow.org, July 2004.