

# How much did shutting down McColo help?

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

## ABSTRACT

On 11<sup>th</sup> November 2008 McColo, a Californian server hosting company was disconnected from the Internet. This took the controllers for several major botnets offline. It has been widely reported that email spam volumes were markedly reduced for some time thereafter. This brief study examines the email arriving at a medium-sized United Kingdom ISP in October and November 2008. It confirms the drop in spam volumes, whilst noting a disproportionate impact on some of the ad hoc measures used to detect email spam.

## 1. INTRODUCTION

On Tuesday 11<sup>th</sup> November 2008 McColo, a Californian server hosting company, was disconnected from the Internet. McColo had a number of customers hosting child sexual abuse images, fake pharmacies, botnet controllers and other wickedness [7], but complaints were proving to be ineffective. Their two Internet connectivity providers, Global Crossing and Hurricane Electric, were shown the results of a Washington Post investigation into McColo's failure to police their customers' activities, and they both decided to withdraw their services [5].

An immediate worldwide drop in email spam occurred, because the command and control systems of six major botnets [1] were no longer in contact with the machines they controlled, preventing any more spam from being sent. Reports at the time showed that spam had dropped to as little as one third of its previous value [8], although the reduction, not unexpectedly, proved to be temporary [6].

In this short paper, I examine the email statistics from a medium sized UK ISP, confirming that a drop in spam volume did occur, but noting some other aspects of the email statistics which show that different peoples' experiences of the event will have had some significant variations.

## 2. WHAT THE DATA SHOWS

The dataset analysed in this paper is for the incoming email to a United Kingdom ISP with *c* 150 000 customers: a mix of individuals, and small and medium-sized businesses. Traffic data was examined for the seven week period 6 Oct–24 Nov 2008. The ISP operates a pipeline of spam detection methods, culminating in a content filtering system provided by Cloudmark. As can be seen in Figure 1, throughout the




Figure 1: Non-spam email volumes were fairly consistent between October (weeks 1–4) and late November 2008 (week 7).

relevant period a fairly consistent 1.1 to 1.6 million non-spam emails arrived most weekdays, with around a million less at the weekends. The only significant variations were on Mondays, which sometimes had a weekend pattern of activity, but some weeks resembled the other weekdays.

The total email spam received during this period is shown in Figure 2 (upper). There is a readily apparent drop in volume starting on the Tuesday of week 6 (McColo was disconnected late in the evening UK time). Before then there were huge variations from day to day – very much the pattern observed in 2007 [2], and indicative of a small number of major senders. Once these senders were eliminated, in weeks 6 and 7, there is much less day-to-day variation.

However, this is only part of the story. The ISP studied, like many others, operates a number of anti-spam heuristics that cause email to be rejected before it is passed to the content filtering system. The specific policies differ from one customer domain to another, with some of the more draconian being applied to “virtual ISPs” (where email service is branded for another organisation). Typical policies are rejection of all email to customers who appear to be under specific attack; refusing to accept email from sites in the SpamHaus Policy Blocklist (PBL) [9]; greylisting [4]; and being picky about some aspects of the SMTP protocol. All these policies mean that substantial numbers of emails never reach the Cloudmark content filters.

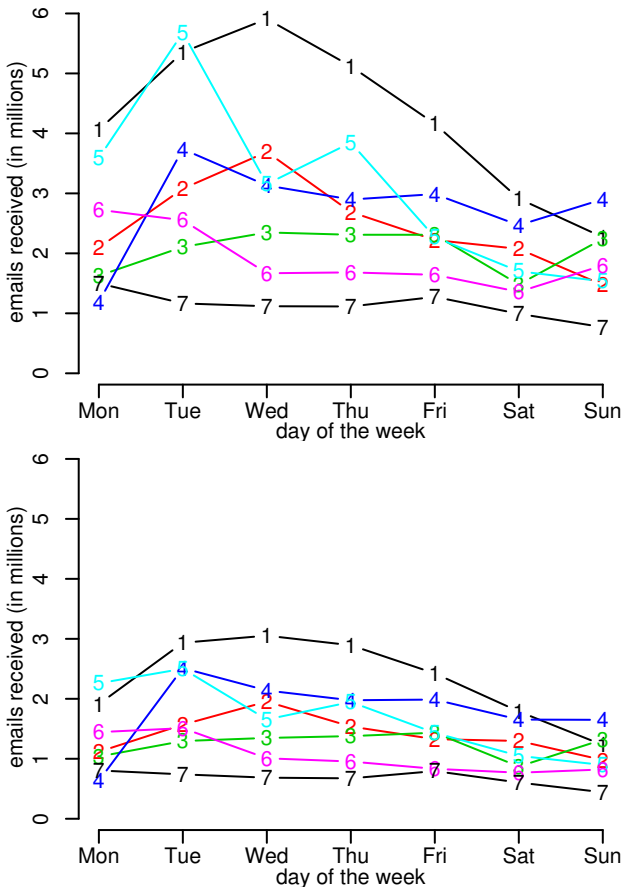


Figure 2: Spam data for October (weeks 1–4) until late November 2008 (week 7). McColo was disconnected late on the Tuesday of week 6. The upper graph is the total spam volume, the lower the spam detected by content filters. Note that after McColo’s disconnection the volumes are more consistent and much less, and that the other content filters are relatively more important.

Figure 2 (lower) plots the volume of spam that was not discarded earlier and hence was only detected by the Cloudmark filters. As can be seen, before McColo’s disconnection between 32% and 56% of the spam need never be assessed by the content filters, thereafter a more consistent 43% or so was detected before the content filter stage.

Figure 3 shows the number of emails rejected because they were sent to non-existent addresses. This mechanism is only employed for a few of the ISP’s customers (hence the lower totals), but on some days it is capable of immediately discarding 900 000 emails. However, once McColo is shut down, a mere 50 000 or so emails a day are blocked.

Finally, Figure 4 considers the subset of email that is discarded because the sender IP address is on a blacklist. This mechanism is only applied to a few of the ISP’s customers, so the actual numbers of emails involved are rather low. This mechanism was being tweaked during the study period, so useful data is only available for the first half of November. Nevertheless, a clear effect is apparent when McColo is shut




Figure 3: Emails rejected because destination did not exist for October (weeks 1–4) until late November 2008 (week 7).




Figure 4: Emails rejected because sender IP address was on a blacklist during the month of November 2008. McColo was disconnected late on the 11<sup>th</sup>.

down. The volume of spam detected this way plummets, falling proportionately more than the overall total. This means that the blacklists ceased to be as useful a spam fighting tool in the immediate aftermath of the McColo closure.

### 3. DISCUSSION

The disconnection of McColo was obviously a Good Thing, because of the substantial, albeit temporary [6], reduction in spam. What the figures presented here have shown is that although overall spam levels fell, particular types of detection mechanism ceased to be as effective. Any system that depended solely on the use of blacklists would have seen a lower percentage reduction than if content filtering was used. Similarly, discarding email to invalid destinations became a far less effective way of reducing the load on spam filtering machines once the botnets were disabled.

As is often the case, people’s experiences of spam will have differed greatly [3]. Headline figures of 60+% reductions only tell one part of a complex story.

## 4. REFERENCES

- [1] N. Chapman. First Atrivo, now McColo, November 2008. <http://www.secureworks.com/research/blog/index.php/2008/11/18/first-atrivo-now-mccolo/>.
- [2] R. Clayton. Email traffic: A quantitative snapshot. In *Proceedings of the Fourth Conference on Email and Anti-Spam (CEAS)*, 2007. <http://www.ceas.cc/2007/papers/paper-76.pdf>.
- [3] R. Clayton. Do zebras get more spam than aardvarks? In *Proceedings of the Fifth Conference on Email and Anti-Spam (CEAS)*, 2008. <http://www.ceas.cc/2008/papers/ceas2008-paper-39.pdf>.
- [4] E. Harris. The next step in the spam control war: Greylisting, 2003. <http://projects.puremagic.com/greylisting/whitepaper.html>.
- [5] J. Hruska. Spam sees big nosedive as rogue ISP McColo knocked offline, November 2008. <http://arstechnica.com/security/news/2008/11/spam-sees-big-nosedive-as-rogue-isp-mccolo-knocked-offline.ars>.
- [6] A. Kleha. Spam data and trends: Q1 2009, March 2009. <http://googleenterprise.blogspot.com/2009/03/spam-data-and-trends-q1-2009.html>.
- [7] B. Krebs. A closer look at McColo, November 2008. [http://voices.washingtonpost.com/securityfix/2008/11/the\\_badness\\_that\\_was\\_mccolo.html](http://voices.washingtonpost.com/securityfix/2008/11/the_badness_that_was_mccolo.html).
- [8] B. Krebs. Spam volumes drop by two-thirds after firm goes offline, November 2008. [http://voices.washingtonpost.com/securityfix/2008/11/spam\\_volumes\\_drop\\_by\\_23\\_after.html](http://voices.washingtonpost.com/securityfix/2008/11/spam_volumes_drop_by_23_after.html).
- [9] The SpamHaus Project. The policy block list, 2009. <http://www.spamhaus.org/pbl/index.lasso>.