

# Spamming Botnets: Are we losing the war?\*

Marios Kokkodis  
University of California Riverside  
mak@cs.ucr.edu

Michalis Faloutsos†  
University of California Riverside  
michalis@cs.ucr.edu

## ABSTRACT

In this work, we examine the spamming activity of the IP space over time, and observe a worrisome phenomenon: spamming botnets are more widely and thinly spread in the IP space over the last four years. We find that (a), a previously-unreported IP space (113.\* - 126.\*) has become a major source of spamming activity, and (b), the spamming activity is more equally distributed among IP addresses. This spreading has grave significance and implications: (1) IP-based filtering for bots and spam will become less effective and more challenging, and (2) system administrators need to improve their mitigation techniques in order to win the war against *botmasters*.

## 1. INTRODUCTION

There is an ongoing battle between *botmasters* and security administrators regarding the proliferation of bots. Bots are compromised machines that are being exploited by a controller (called *botmaster*), for various types of malicious activities. One common activity of a botnet (group of bots) is to launch spam campaigns. Knowing the spatio-temporal behavior of spammers can help us devise efficient spam filters and mitigation techniques [1].

Several studies have examined the distribution of spam activity across the *IP* address space. Specifically, we know the following:

- The majority of spam emails comes from bots [2, 3, 4].
- Two *IP*-spaces are responsible for the majority of the observed spam [3, 5].
- The spam activity seems quite “concentrated” and follows the Pareto principle (the 80-20 rule) [5, 6].

These findings are implicitly optimistic, as they suggest that by focusing on few highly active *IP*-prefixes, we may be able to mitigate the effect of spam activity.

In this work, we study spamming botnets over the span of 4 years (2006 - 2009), using a new user-collected dataset. This dataset, consists of 2, 046, 520 different spam emails,

\*This work was supported partly by NSF CT-ISG 0831530.

†Also affiliated with *StopTheHacker.com*

collected by user accounts from three different domains, between 01 - 01 - 2006 and 05 - 31 - 2009. The data analysis unveils a worrisome phenomenon: bots seem to have spread widely across the *IP* space. Our contributions can be summarized in the following:

- We identify a **new** *IP*-space area with very high spamming activity, in the range between 113.\* and 126.\*, which first appears in 2007.
- We observe that the appearance of spamming bots becomes more thinly spread.


These two points suggest that *IP*-based filtering is becoming more challenging. In addition, it seems that our efforts to contain and eradicate bots may *not* be as successful as we would like them to be.

## 2. DATA COLLECTION

For our analysis, we use a new publicly available dataset [7] which consists of 2,046,520 spam emails. Those messages were collected by various user accounts from three different domains, over the span of four years: 2006, 2007, 2008 and 2009. The dataset of each year consists of 699K, 324K, 725K and 297K spam emails, respectively<sup>1</sup>. Note that the data for 2009 has not yet been completed at the time of the study. The majority of these emails were flagged as spam by *Spam-Assassin* [8], a well known email filtering application. To increase our confidence in the dataset, we manually verified a randomly chosen subset of the emails.

For each email in the dataset we try to find its source *IP*. According to the SMTP protocol [9], each server that receives a message, appends a *Received* record (e.g. *Received: from example.com ([77.49.119.108])*) to the top of the email header. Hence, the earliest *Received* record should include the *IP* of the first SMTP server that forwarded the email (i.e. the source). However, in the case of spam, the protocol is often violated, since spammers have developed techniques to obfuscate their identities. An example of such a technique is to falsify the header information by either modifying it or by appending invalid *Received* headers [10]). Therefore, the only relay which we can identify the true *IP* address, is the one that established the SMTP connection to our mail server. In the following, we use these *IP*s for conducting our analysis. Note that the same method is used in [3, 11, 12].

<sup>1</sup>The reduction in the amount of spam from 2007 to 2008 is due to the shut down of a spam trap. In 2008, we receive more spam than 2006 and 2007, without the use of the trap.



**Figure 1: The distribution of spamming activity across the IP space over the last four years. We show the two highly active spamming areas (left and right boxes) and an emerging high activity area (middle box) not reported so far.**

### 3. SPAM DISTRIBUTION OVER TIME

Next we describe the conducted analysis along with the derived results. The first part provides some statistics on our data, while the second and the third one discuss our findings.

#### 3.1 Data Analysis

In Figure 1, we present the CDF of the spammers’ IP addresses, for all four datasets.


For 2006, there are three high activity spamming chunks of IPs (presented in Table 1, spaces A and C), which constitute 22.6% of the total IP address space, and are also the origin of 92% of the total amount of spam that we receive in 2006. In comparison to these findings, the corresponding 2007 high-activity areas (presented in the second row of Table 1) cover 29.3% of the total IP space, and are responsible for 95% of the total volume of received spam. In addition, the high-activity chunks of 2006, are only a subset of the respective spamming chunks of 2007, an observation that shows a spread of spammers over the IP space (discussed in §3.3)

In 2008, we identify three high activity areas (third row in Table 1). These are the cause of 91.5% of the total volume of spam, and constitute 32.4% of the total IP address space. A similar argument can be made for the high-activity spamming IPs of 2009 (fourth row in Table 1). These chunks are responsible for 93.4% of the total amount of spam, while they cover 34.4% of the IP space.

#### 3.2 New Actively Spamming IP Space

**FINDING 1.** *We observe a previously-unreported IP space (113.\* - 126.\*) with intense spamming activity in our dataset.*

Figure 1 shows the CDF of the spam distribution per year. In Table 1, we list the highly active IP-spaces regarding spamming activity in our dataset. To the best of our knowledge, no one so far had observed high spam activity in IP space B: between 113.\* and 126.\* (shown as an inset in Figure 1). Note that spaces A and C in the table were reported



**Figure 2: The cumulative percentage of spamming contribution of the common /16 prefixes for the last four years. The activity becomes less concentrated from 2006 to 2009, which indicates that more spam is distributed among new /16 subnets across the IP space.**

by previous studies [3], which increases the confidence of our dataset. This new space of highly active IPs shows low spam intensity in 2007, and by 2009, it becomes one of the three major spamming IP-areas: it is the origin of 15% of the received volume of spam in 2009 while it constitutes only 5% of the total IP space.

Year	space A	Space B (new)	Space C
2006	58.* - 73.* 80.* - 90.*	-	190.* - 222.*
2007	57.* - 92.*	121.* - 126.*	188.* - 222.*
2008	57.* - 95.*	116.* - 126.*	188.* - 222.*
2009	57.* - 97.*	113.* - 126.*	188.* - 222.*

**Table 1: High Activity Spamming sets of IPs**

#### 3.3 The Spreading of Spam Activity

**FINDING 2.** *We observe a “spreading” trend of the spamming activity in the IP space from 2006 to 2009 in our data.*

From the analysis we conducted in §3.1, we observe that spamming activity is “spreading” over the IP space in two ways: (a) a new active area emerges (2007-2009, §3.2), and (b) the known major spamming areas become wider, as shown in table 1 and figure 2.

There are several different ways to quantify this “spreading”. For example in 2006, the highly active spaces cover 22.6% of the total IP space. This percentage increases every year and reaches 34.4% in 2009 (discussed also in §3.1), illustrating the trend of the spamming areas to spread across the IP space.

Another way to show this “spreading”, is to focus on the spam activity of the /16 subnetworks that were active in all our datasets. In Figure 2, we plot the cumulative percentage of the spam activity of these /16 prefixes, as a percentage of the total received spam of each year. Note that the total on the y-axis does not sum up to 100%, as there is contribution from /16 subnets that are not part of the group we examine. The x-axis presents the active /16 prefixes, in order

of decreasing activity. Intuitively, the closer the line is to the upper left corner, the more concentrated the spamming activity. In 2006 almost 90% of the total volume of received spam originated from those /16 subnets. In the following years, the contribution of these subnets steadily decreases, dropping down to 52% in 2009. This indicates that through time, new *IPs* are responsible for an increasing portion of the total volume of received spam.

## 4. CONCLUSION

In this work we analysed a new publicly available spam archive which consists of spam gathered during the last 4 years. Our analysis revealed a worrisome observation: spamming bots seem to spread wider across the *IP* space since 2006. This spreading has two major implications: (a) *IP*-based filtering for bots and spam is becoming less effective and more challenging, and (b) security administrators may be losing the war against *botmasters*. In the future, we intend to further document this spreading, and also correlate it with the locations and networks that bots seem to be appearing in.

## 5. REFERENCES

- [1] Anirudh Ramachandran, David Dagon, and Nick Feamster. Can dns-based blacklists keep up with bots? *CEAS*, 2006.
- [2] Husain Husna, Santi Phithakkitnukoon, Srikanth Palla, and Ram Dantu. Behavior analysis of spam botnets. *COMSWARE*, 2008.
- [3] Anirudh Ramachandran and Nick Feamster. Understanding the network-level behavior of spammers. *SIGCOMM*, 2006.
- [4] Yinglian Xie, Fang Yu, Kannan Achan, Rina Panigrahy, Geoff Hulten, and Ivan Osipkov. Spamming botnets: Signatures and characteristics. *SIGCOMM*, 2008.
- [5] Zesheng Chen, Chuanyi Ji, and Paul Barford. Spatial-temporal characteristics of internet malicious sources. *INFOCOM*, 2008.
- [6] Wikipedia. [http://en.wikipedia.org/wiki/pareto\\_principle](http://en.wikipedia.org/wiki/pareto_principle).
- [7] Untroubled. <http://untroubled.org/spam/>.
- [8] SpamAssassin. <http://spamassassin.apache.org/>.
- [9] Simple Mail Transfer Protocol. <http://www.ietf.org/rfc/rfc2821.txt>.
- [10] Joshua Goodman. Ip addresses in email clients. *CEAS*, 2004.
- [11] John P. John, Alexander Moshchuck, Steven D. Gribble, and Arvind Krishnamurthy. Studying spamming botnets using botlab. *USENIX NSDI*, 2009.
- [12] Yinglian Xie, Fank Yu, Kannan Achan, Eliot Gillum, Moises Goldszmidt, and Ted Wobber. How dynamic are ip addresses. *SIGCOMM*, 2007.