Understanding Large-Scale Spamming Botnets From Internet Edge Sites

Tatsuya Mori
NTT Laboratories
3-9-11 Midoricho Musashino
Tokyo, Japan 180-8585
mori.tatsuya@lab.ntt.co.jp

Holly Esquivel
UW - Madison
1210 W. Dayton St.
Madison, WI 53706-1685
esquivel@cs.wisc.edu

Aditya Akella
UW - Madison
1210 W. Dayton St.
Madison, WI 53706-1685
akella@cs.wisc.edu

Akihiro Shimoda
Waseda University
3-4-1 Ohkubo, Shinjuku
Tokyo, Japan 169-8555
shimo@goto.info.waseda.ac.jp

Shigeki Goto
Waseda University
3-4-1 Ohkubo, Shinjuku
Tokyo, Japan 169-8555
goto@goto.info.waseda.ac.jp

ABSTRACT

This paper aims to understand empirically the impact of a large-scale spamming botnet, and the effectiveness of targeting its core infrastructure – C&C servers – from the viewpoint of several Internet edge sites. We also attempt to study the characteristics of the spamming botnet in the long-term to see how quickly bot masters react and what type of action they take. Our primary target in this paper is one of the world's previously worst known spamming botnets, Srizbi, whose C&C servers were shutdown by its upstream ISPs on November 11, 2008. We conduct an extensive measurement study spanning a large volume of e-mail delivery logs and packet traces collected at five vantage points. The measurement period spans three years and includes of the rise and fall of the botnet. We leverage passive TCP fingerprinting on the collected packet traces to identify bot-infected hosts and spam messages sent from them. We first extract variants of the known TCP signatures that are associated with the spamming botnet by correlating the data sets in the time and space domains. Next, by using the signatures, we quantify the volume of spam sent from the botnet and the effectiveness of the C&C server shutdown from an Internet edge site-perspective. We attempt to study the characteristics of the spamming botnet in both the time and space domains. We reveal several findings that are useful in understanding the spread of spamming botnets; specifically, we note the steady growth of the botnet's size and the rapid version transition after the shutdown of C&C servers. We also estimate the entire size of Srizbi botnet. We then study how the botnet membership is distributed around the globe and how its size changed over time.

1. INTRODUCTION

Over the past few years, the volume of e-mail spam has grown significantly to the point it is no longer just a nuisance. Some reports suggest that as much as 90–95% of all e-mail sent or received today is spam [1, 6]. Today, botnets are widely used as a *scalable* and *elusive* approach to disseminating spam messages. Spammers purchase access to a fraction of bots controlled by the botnet to send out spam messages from the infected hosts. Spammers send these instructions from the command and control (C&C) server via encrypted channels to the bots. Recently, spamming botnets have made the transition from proxy-based spamming to template-based spamming. These new sophisticated user interfaces play a key role in the efficiency of dissemination mechanisms in spamming infrastructures [9]. These improvements have lead to an exponential increase in spamming capabilities. For example, "Srizbi" is claimed to be capable of sending 60 billion spam messages per day, which is more than half of the total 100 billion spam messages sent per day on average [10]. Another large-scale spamming botnet, "Conficker", consists of more than 10 million infected hosts all over the world and could be capable of sending out 400 billion spam messages per day [3]. These large global-spamming infrastructures have traditionally been hard to stifle.

In late 2008, a bold and drastic action was taken to contain the world's worst spamming botnet, Srizbi. On November 11, the web hosting service provider, McColo, was shut down by its two upstream ISPs. McColo is known as a so-called "bulletproof hosting" company because it reportedly allowed its customers to bypass laws regulating Internet content and services. McColo also reportedly allowed these customers to remain online regardless of complaints. The company hosted the C&C servers for major spamming botnets, including Srizbi [4]. Accordingly, as many operators and researchers expected, it is widely reported that the drop in spam volume was estimated to be between 50 to 75 percent on the very same day [4, 15].

Although some measurement studies reported that spam volume have returned to pre-McColo shutdown levels [15], the temporary but great success of the shutdown indicates that this unprecedented and drastic move was effective. This action allows us to better understand the larger picture of spamming botnets and the way in which they can make transitions, which is crucial to building effective and sustainable anti-spam solutions. As a first step toward this goal, we aim to understand the world's worst spamming botnet, *Srizbi*, and to study the effectiveness of targeting the botnet's C&C servers (i.e., McColo shutdown). We also look at the long-term trends of Srizbi to study how the botnet has grown and reacted to the shutdown.

We conduct an extensive analysis of e-mail delivery logs and packet traces (tcpdump) collected at five different vantage points across two countries: US and Japan. We also use publicly available packet traces published by MAWI [13]. The five locations

CEAS 2010 - Seventh annual Collaboration, Electronic messaging, Anti-Abuse and Spam Conference July 13-14, 2010, Redmond, Washington, US

consist of four different types of Internet edge sites, namely, an enterprise network, a campus network, a leaf site of a scientific research network, and an international backbone link used by several research organizations. The total data collection periods span from July 2007 to November 2009.

To detect Srizbi bots, we leverage TCP fingerprinting, which can identify the operating system of a host based on the TCP/IP stack of the system. As Stern discovered [20], Srizbi uses a dedicated network driver that uses intrinsic TCP/IP parameter settings. Thus, we can extract hosts infected with the Srizbi trojan by tracking their TCP fingerprint signature. In addition, we conduct temporal and spatial analysis of spam sending patterns of bots with specific other signatures to extract variants of the basic set of Srizbi signatures that were identified in [20]. Finally, we correlate e-mail delivery logs with packet traces and identify the volume of spam sent by Srizbi botnets using the extracted signatures.

The primary contributions of this work are:

- We extract new previously unknown variants of TCP signatures that are also associated with the botnet to aid in the detection of bots and spam sent from them.
- We quantify the volume of spam sent from the botnet and study the effectiveness of the shutdown of its C&C servers from the view point of Internet edge sites.
- We reveal several findings that are useful in understanding the spread of spamming botnets; specifically, we note the steady growth of the botnet's size and the rapid version transition after the shutdown of C&C servers. To the best of our knowledge, our work is the first one that clearly reveals this drastic version transition.

We argue that the analysis of long-term data sets collected at multiple vantage points, i.e., Internet edge sites, can help in understanding how quickly a botnet could grow, how long the botnet could stay active, how large the botnet could become, how they could be mitigated by an action against their core infrastructure, and how quickly they could recover from the action. These observations are essential as a first step toward building a method to inactivate spamming botnets permanently.

The remainder of this paper is structured as follows: Section 2 presents a description of data sets utilized in this work. In Section 3, we present our findings on the characteristics and trends of the Srizbi botnet. In Section 4 we discuss related studies and how they compare to ours. Finally, Section 5 concludes our work.

2. DATA DESCRIPTION

We collected data from five vantage points located at different organizations and countries. The measurement period spans two years, from July 2007 to November 2009. The data sets were collected at the University of Wisconsin - Madison, USA; Waseda University in Japan; a middle size corporation in Tokyo, Japan; a leaf site of the scientific research network, GEMnet2 [22]; we also use publicly available data published by the MAWI WG of the WIDE project [13]. In this work, we call these vantage points UW, WAS, CORP, GEM, and MAWI, respectively.

Each vantage point collects one or two primary data sets that are used for this spam analysis. The first set of data collected consists of packet traces of all incoming TCP SYN packets to the SMTP (Simple Mail Transfer Protocol) servers collected using tcpdump. The second set of data contains all SMTP delivery records for each vantage point for the respective e-mail servers. Table 1 summarizes the measurement periods of each data set. In the following, we describe the data sets in detail.

Table 1: Total measurement period for each data set.

	tcpdump	SMTP log
UW	Feb 9, 2008 – Jul 11, 2008	Feb 1, 2008 – Apr 30, 2008
CORP	Apr 7, 2008 – Jul 31, 2008 Dec 26, 2008 – Dec 31, 2009	Apr 7, 2008 – Jul 31, 2008 Jan 1, 2009 – Dec 31, 2009
GEM	–	Aug 1, 2008 – Apr 30, 2009
WAS	Oct 16–22, 2008, Mar 7–16, 2009	–
MAWI	Jul 1, 2007 – Nov 31, 2009	–

Table 2: Statistics of the SMTP logs for selected months.

	#spam	#ham	#senders
Pre-McColo			
UW Apr 2008	101,131,663	12,265,296	7,473,847
CORP Apr 2008	20,107,288	545,686	2,590,289
GEM Aug 2008	95,405	1,067	68,100
Post-McColo			
CORP Jan 2009	10,886,153	723,142	1,236,965
GEM Dec 2008	65,491	2,588	36,344

2.1 Tcpdump

For UW, WAS, and CORP, packet traces are collected on the incoming external links of the networks. For MAWI, we use packet traces which were collected on trans-Pacific line (150-Mbps link) that connects US and Japan, which is utilized by several research organizations. Analyzing packet traces enables us to study all the incoming SMTP connections to the networks. To extract minimal information, we filter all the packets other than TCP packets with SYN flag that are destined to the SMTP port. This allows us to employ TCP fingerprinting on packets from e-mail senders while discarding all other private information in the subsequent SMTP transmissions. Since MAWI traces have been collected since July 2007, we can study the long-term trends of the spamming botnets.

2.2 SMTP logs

For UW, CORP and GEM, SMTP logs were collected on commercial anti-spam appliances that work as MX servers. UW and CORP operate greylisting mechanisms at the MX servers. Greylisting is a mechanism that temporarily rejects e-mail messages from a sender which has not previously been seen. Greylisting is effective because if an e-mail is rejected, a spammer will likely not retransmit it since spammers cannot afford the time and resources to retry thousands of bounced messages. By analyzing greylisting logs, the SMTP connections which did not attempt retransmission can be extracted. In this work, we regard these connections as *attempted* spam messages sent to the e-mail servers. That is, if a connection is filtered by greylisting and is not retried later, we regard the connection as a spam message. Note that most spam messages were filtered at the greylisting stage in our data sets. The anti-spam appliances then apply content-based filtering to all messages which *pass* the greylist filtering and spam scores are assigned to them. We adopt conservative thresholds to classify e-mail messages into spam, or ham, based on the score. For example, a spam e-mail must have a spam probability score of greater than 0.95 out of 1.0 in order to be considered spam, while a ham or legitimate e-mail must have a score of smaller than 0.05. In the data sets we analyzed, the majority of messages and connections are classified into spam or ham with the definitions shown above. We note that software-based

filtering is error-prone and thus could affect the derived statistics, but we expect our high-level observations to remain qualitatively similar.

Table 2 shows the resulting classification statistics of the logs for selected months. We note that majority of messages seen in all data sets is spam, which is consistent with previous observations [1].

3. ANALYSIS

In this paper, we aim to understand the world's worst spamming botnet, *Srizbi* and to study the effectiveness of targeting the botnet's C&C servers, i.e., *McColo shutdown*. We also attempt to study the characteristics of the spamming botnet in both time domain and space domain. First, we show how we identify the infected-hosts using TCP fingerprinting. We present new variants of TCP signatures that are also associated with the spamming botnet. (Section 3.2). Second, we quantify the volume of spam sent from the Srizbi botnet, and study the effectiveness of the shutdown from the view point of Internet edge sites (Section 3.3). Third, we reveal the growth of Srizbi botnet and the version transition of Srizbi around the shutdown period (Section 3.4). Finally, we attempt to characterize the Srizbi botnet. We first estimate the entire size of Srizbi botnet and how it changes over time. We also look at the global correlation of the botnet activities observed in two different countries (Section 3.5.1). Next, we study how the infected bots were distributed over the global Internet (Section 3.5.2).

3.1 Overview of approach

The key to our work is analyzing long-term data sets collected at Internet edge-sites to extract useful information for understanding botnets. We utilize TCP fingerprinting to identify the operating system characteristics of hosts infected with the Srizbi bot. The signatures are extracted by employing p0f [24] over the collected tcpdump files. The format of the extracted signature is

- `[W:T:D:S:O...:Q]`

where `W` stores the information about the window size, `T` is the initial value of Time to live (TTL) field, `D` is the do-not fragment (DF) bit, `S` is overall SYN packet size, `O` is the option value and order specification, and `Q` is a list of miscellaneous information. A full description of option values and miscellaneous information can be found in [24]. This process helps identify bots based on known signatures, as well as, extract variant signatures.

In order to associate spam messages and their respective TCP fingerprints, the tcpdump and SMTP logs are correlated together based on timestamps. We note that the join of SMTP and tcpdump logs is 1-to-1. All spam messages that appear in the SMTP logs are mapped back to their associated TCP fingerprints found in the tcpdump logs. This is done with the UW and CORP data sets to help identify variant signatures, calculate spam statistics, and evaluate the effectiveness of targeting the core infrastructure of spamming botnets.

Next, we use the MAWI data set to study long-term trends. MAWI captures packet traces for 15 minutes each day from from 14:00 to 14:15. Although the measured information is sampled in time (sampling rate is 1/96), the data set is useful to track long-term trends and activity of the botnet, and identify potential signature variants.

3.2 Identifying spamming botnets

The first step in understanding a botnet is identifying the infected hosts which belong to it with a high degree of accuracy. Stern [20] carefully studied hosts infected with the Srizbi trojan and found that Srizbi's TCP/IP driver uses a rare combination of parameters, which are not used by other operating systems listed in the p0f signatures[1]. According to Stern [20], there are two sets of Srizbi signatures, one prevalent before the McColo shutdown and the other afterward. We refer to these sets of signatures as Srizbi V1 and V2, respectively.

Table 3: Top 5 spam-sending signatures of Srizbi V1 (known are in bold font) and their potential variants for UW (top) and CORP (bottom) in April, 2008.

signature	#spam	#ham	#senders
UW			
[24000:128:0:44:M536:.]	14,495,869	2,708	260,955
[24000:128:0:44:M1360:.]	262,077	21	3,147
[24000:128:0:44:M528:.]	223,246	3	2,662
[24000:128:0:44:M1452:.]	56,589	9	774
[24000:128:0:44:M1414:.]	20,504	7	251
CORP			
[24000:128:0:44:M536:.]	7,252,084	41	1,139,778
[24000:128:0:44:M1360:.]	126,955	0	21,329
[24000:128:0:44:M528:.]	90,518	0	9,463
[24000:128:0:44:M1452:.]	30,660	0	4,025
[24000:128:0:44:M1414:.]	12,109	0	2,428

The following are the three *known* botnet signatures:

- `[24000:128:0:44:M536:.]` (Srizbi V1, Ethernet)
- `[24000:128:0:44:M528:.]` (Srizbi V1, ADSL)
- `[6144:255:0:44:M1024:.]` (Srizbi V2)

In addition to the above signatures, we found several variants that are associated with Srizbi (V1 and V2) with high probability. We identified these variants as follows: First, we note that the fraction of spam messages sent by the Srizbi signatures are quite high. Thus, we identify signatures that exhibit a similarly high fraction of spam messages and similar TCP characteristics. Table 3 shows the top 3 similar spam sending signatures (plus two known signatures in bold). Interestingly, the top 5 signatures and their ranking were in common among the UW and CORP data sets. These results hold for later months which indicates that these signatures were stable over time (before shutdown).

These variants only differ in their MSS values, which reflects the varying window sizes imposed by different types of Internet access links. Table 3 shows the top 5 spam-sending signatures for Srizbi V1 and their variants, sorted by total number of spam messages in the two data sets, UW and CORP. Similarly, Table 4 shows the top 5 packet-sending signatures of Srizbi V1 and variants in a week packet trace of WASEDA. Here, the value of initial TTL, T, is corrected with the formula $T = 2^{\lceil \log_2(t)+1 \rceil}$, where t is the value of observed TTL. For instance, we observe $t = 49$, it is corrected to its potential initial value, $T = 64$ [2].

Next, as we detail in Section 3.4, the long-term history of these variant signatures almost exactly agrees with the original ones. Unlike the previously known signatures, these signatures send spam on much smaller scale, but the same rise and fall patterns in spam volume can be observed before and after the shutdown.

[1]We manually collected newer operating signatures that are not listed on the original p0f signature list, e.g., Windows Vista and Mac OS X 10.5+, and found that none of them matched the extracted Sribzi signatures.

[2]If the corrected value is 256, we apply another heuristics, i.e., the value is corrected as 255. Note that if hop length between sender and measurement point is too long, say, more than 32, the estimated TTL can be smaller than actual value.

Table 4: Top 5 signatures of Srizbi V1 (known signatures in bold font) and variants for WAS in 16–22 Oct, 2008.

signature	#pkts	#senders
[T16:128:0:44:M536:.]	1,027,408	160,760
[T16:128:0:44:M1360:.]	20,113	3,079
[T16:128:0:44:M528:.]	6,269	1,102
[T16:128:0:44:M1452:.]	5,339	661
[T16:128:0:44:M1400:.]	1,679	112

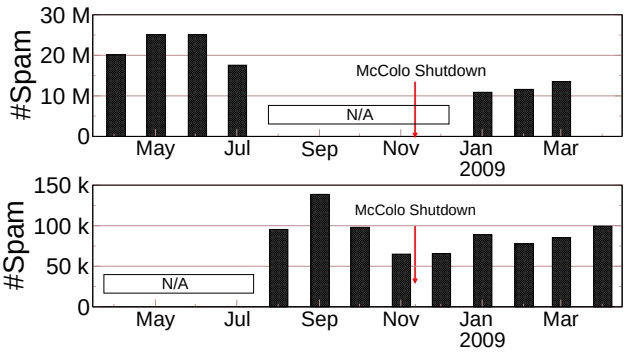

Figure 1: The volume of spam for CORP (top) and GEM (bottom) data sets observed by month .

Finally we note, that the above described patterns hold for Srizbi V2 signatures. We omit these results for brevity. Thus, we conjecture that these variants are also associated with the Srizbi botnet. Based on these observations, we add the following signatures as variants of the known Srizbi V1/V2 signatures.

- `[24000:128:0:44:M*:.]` (Potential Srizbi V1 variants)
- `[6144:255:0:44:M*:.]` (Potential Srizbi V2 variants)

In the following sections, we leverage these signatures to study the scale of Srizbi botnet, as well as its impact and long-term growth and evolution.

3.3 Effectiveness of targeting C&C servers

Here, we study how the volume of spam changed after the shutdown of C&C servers from the view point of Internet edge sites. Figure 1 shows the total received spam volumes for CORP and GEM over a period of several months. In both cases, we can see a large reduction in spam volume after the shutdown. However, in April 2009, the spam volume for the GEM data set has almost returned to the pre-shutdown level, which agrees with the observation in [15]. On the other hand, the spam volume for the CORP data set has remained at a lower level, i.e., about half of the peak volume, for more than 4 months after the shutdown. According to the network operator of CORP, the spam volume was still lower as of May 2009—6+ months after the shutdown. Thus, the long-term effectiveness of the shutdown varies at Internet edge sites.

Next, we study the decrease in spam volume observed from Srizbi botsas a result of the shutdown of its C&C servers. Based on the identification techniques described in Section 3.2 and by correlating the two sets of data we have (SMTP and tcpdump), we identify spam messages sent by Srizbi bots. Table 5 shows the numbers and fractions of spam messages attributed to Srizbi hosts at edge-sites

Table 5: Breakdown of spam messages sent from end-hosts with Srizbi or Windows-based TCP fingerprint signatures.

data set	#total spam	Srizbi (%)	Windows (%)
Before shutdown			
UW Feb 2008	110,959,667	12,602,852 (11%)	83,333,645 (61%)
UW Mar 2008	136,572,281	17,813,844 (13%)	101,094,771 (74%)
UW Apr 2008	101,131,663	15,185,849 (15%)	71,106,454 (70%)
CORP Apr 2008	20,107,288	7,530,864 (37%)	11,220,937 (56%)
CORP May 2008	25,079,293	10,694,254 (43%)	13,286,069 (53%)
CORP Jun 2008	25,088,872	11,349,148 (45%)	12,707,436 (51%)
CORP Jul 2008	17,562,162	5,434,277 (30%)	10,682,847 (60%)
After shutdown			
CORP Jan 2009	10,886,153	607,499 (6%)	9,487,679 (87%)
CORP Feb 2009	11,604,039	951,914 (8%)	9,849,693 (85%)
CORP Mar 2009	13,545,628	246,862 (2%)	12,211,121 (90%)

UW and CORP. Prior to the shutdown, a large number of spam messages were sent by Srizbi, but the fraction that Srizbi contributed to total spam volume differed from site to site. For UW, the fraction of Srizbi spam is around 11–15%. On the other hand, for CORP, the fraction is around 30–45%. We conjecture that the difference in the number spam messages reflects the way in which the recipients' e-mail addresses are harvested by spammers.

We also analyze the source of the remaining spam messages. Table 5 shows the fraction of spam messages sent from Windows hosts. Although the percentage of Windows-based spammers pre-shutdown is lower than reported by previous studies, the post-shutdown fractions are similar to those seen by Ramachandran and Feamster [17]. We check the IP addresses of these hosts against commercial DNS-based Blackhole List (DNSBL) (spamhaus PBL [21]). We found that roughly 90% of IP address belonged to the dynamic IP address space. Because of the large number of end-user machines that belong to this range, we conjecture that throughout the entire measurement period a large fraction of these hosts were likely infected with bots (perhaps non-Srizbi). This indicates that potentially targeting other C&C servers could reduce spam volumes even further.

After the shutdown, we see a significant reduction in spam volumes, especially those for Srizbi V1. The total volume of observed spam messages for the CORP data set is reduced to roughly 50% of the pre-shutdown level. This indicates that the shutdown effectively reduced the number of spam messages seen, and hindered a previously prevalent global-scale spam sending infrastructure. Although spam volumes rebounded, continuing action, such as the shutdown of other source controllers, could be debilitating to some botnet infrastructures.

3.4 Long-term analysis

This section analyzes the long-term trends of the Srizbi botnet including its rise, fall and version transition. Figure 2 shows the number of observed IP addresses with Srizbi V1/V2 signatures and their variants in the MAWI data set. The first packet from Srizbi V1 was observed on August 7, 2007. The number of Srizbi hosts observed exceeded a hundred two weeks later and the number kept growing steadily as depicted in the figure. On November 11, 2008 we see the fall of Srizbi V1 on the day of the McColo shutdown.

Interestingly, we show Srizbi V2 and its variants have been active since late October 2008. These bots were soon terminated with the same shutdown that effected V1. The activity of Srizbi rebounded about two weeks later but this time, only V2 and its variants survived (see Figure 2). According to studies such as [2], the malware leveraged a *rapid* fallback mechanism to update itself, which included references to a new set of C&C servers. To the best of our

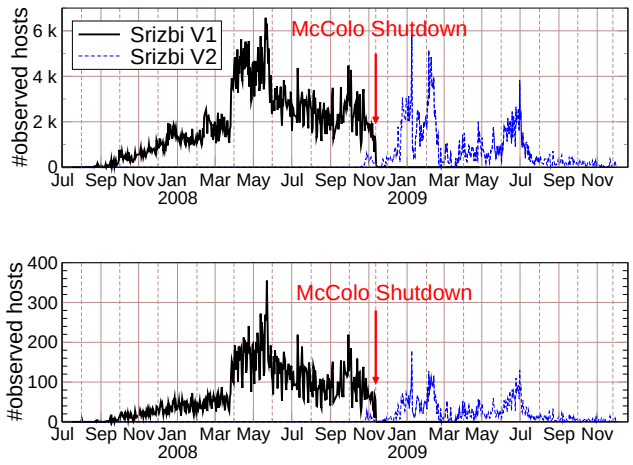

Figure 2: Number of observed hosts infected with Srizbi V1/V2 (top) and their variants (bottom) in the MAWI data set.

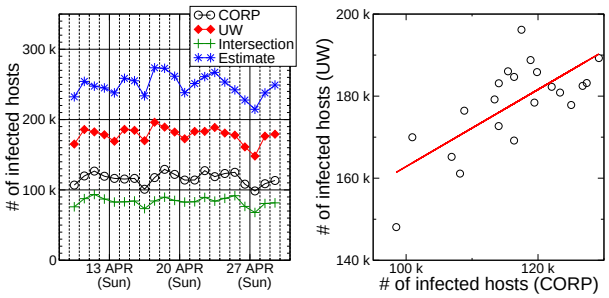

Figure 3: (Left) Estimation of the size of the Srizbi botnet in April 2008. (Right) Scatter plot of active bot sizes observed at CORP and UW in April 2008. The line indicates the linear regression.

knowledge, while many studies such as [15, 20] have reported the resurrection of Srizbi after the shutdown, our study is the first one that quantitatively presents this rapid version transition around that time period. After the transition, we notice that Srizbi V2 has been less widely spread, compared to V1 before the shutdown. We observed similar results in other data sets.

As Stern indicated in [20], in February 2009 a signature for Srizbi was added to the Malicious Software Removal Tool [14], thus potentially mitigating the spread of V2. Unlike V1 though, the new version of the malware has still been active months after the update (see Figure 2). Although we see some fluctuations in mid 2009, the Srizbi botnet has been minimally active since November 2009 – a year after the shutdown. We note that a few e-mails from V2 are still seen as of April 2010, but the number is exponentially smaller than V2 at its peak in July 2009.

3.5 Characterizing Srizbi

3.5.1 Estimating Size of Srizbi

Knowing the scale of spamming botnet is useful to estimate the possible worst-case damage caused by a spam flood from a botnet. We leverage a technique proposed by Lawrence and Giles [11] to estimate the size of the Sribzi botnet in a probabilistic way. They estimate the size of indexable web pages on the Internet through the analysis of collected web pages by search engines. To do this, they leverage independently sampled data.

Let $P(X)$ be the probability that a spam bot hits the vantage point X. If we assume that two vantage points A and B receive spam messages from the Srizbi botnet independently, i.e., a bot selects recipients of spam messages randomly, the probability that a spam bot hits both vantage points is given as $P(A, B) = P(A)P(B)$. Therefore, the total number of hosts infected with Srizbi, $N(\Omega)$, can be estimated as $\widehat{N(\Omega)} = N(A)N(B)/N(A, B)$, where Ω is the entire Internet space and $N(X)$ is the number of spam bots that hit the vantage point X. In this analysis, we use the tcpdump logs of UW and CORP data sets. Because of differing types of these organizations, it is natural to assume that the botnet hits these two sites independently.

When estimating the number of infected hosts it is necessary to take account the reassignment of IP addresses. Zhuang et al. studied the dynamics of IP addresses through the extensive analysis of user login/logout events on Hotmail [25]. They found that about 25% of IP addresses never see IP reassignment in a 7 day period, while a large portion of IP addresses get reassigned almost every day. Based on the above observations, we assume that majority of IP addresses assigned to hosts are stable on a given day; thus, the number of infected hosts seen on a day can be estimated by counting the number of distinct IP addresses seen on that day.

The left panel of Figure 3 shows the number of IP addresses per day for each data set, their intersections, and the estimated number of active Srizbi-infected hosts per day using the probabilistic model. The analysis uses the data sets collected from 00:00:00, April 9, 2008 to 23:59:59, April 29, 2008 in UTC timezone. We note that the offset of timezones for both sites are corrected. The estimated values range from 210K hosts per day to 275K hosts per day. These numbers agree with the other estimates previously reported in [20] and [9], which claimed that the lower bound of Srizbi botnet size is around 80-130K per day in April 2008 [20] and, the size of Srizbi botnet was around 315K hosts per day in April 2008 [9], respectively.

We also notice that there is clear time synchronization between the number of infected hosts observed at each location. The right panel of Figure 3 shows a scatter plot of this trend. We see positive correlation between them with a resulting correlation coefficient of 0.715. We conjecture that the time synchronization reflects the activity of the end-hosts. For instance, the number of hosts decreases every Sunday (in UTC). The way in which a botnet is used, e.g., size of spam campaigns, may also contribute to the global synchronization of botnet activity and the effect of the C&C server shutdown.

In contrast, Figure 4 shows the same analysis for the data set observed in November 2009. The estimated size ranged from 1.5K to 6.2K hosts per day. The size of the botnet has decreased by two orders of magnitude. However, we note that a considerable number of hosts are still active in the Internet. In Figure 4, we cannot see the time synchronization between the infected hosts for each location like previously. The correlation coefficient between them is 0.003; which means the botnet activity is no longer synchronized. We conjecture that spammers are not continuously using the botnet to send out spam like before.

3.5.2 Spread of Srizbi

Finally, we study the spread of Srizbi-infected hosts around the

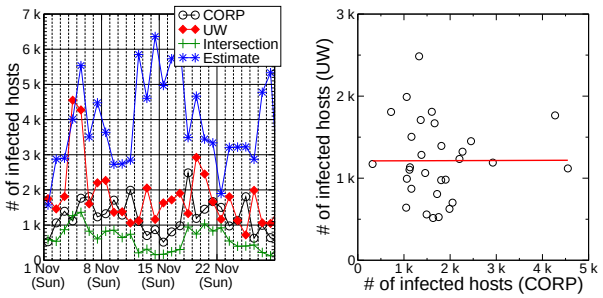

Figure 4: (Left) Estimation of the size of the Srizbi botnet in November 2009. (Right) Scatter plot of active bot sizes observed at CORP and UW in November 2009. The line indicates the linear regression.

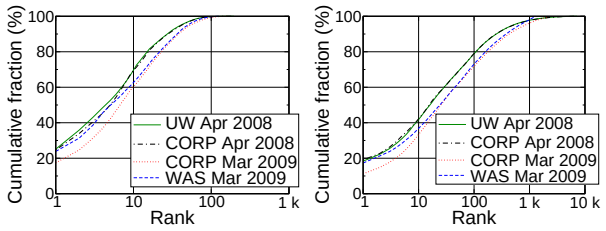

Figure 5: Cumulative fraction of the Srizbi hosts in top countries (left) and ASNs (right).

world. Tables 6 and 7 show the top 5 countries and ASNs of Srizbi-infected hosts before and after the McColo shutdown. Before the McColo shutdown, 4 out of 5 countries, namely Turkey, Russia, Brazil, and US, were in common among the three vantage points. More than 40% of Srizbi hosts belong to these top countries. This observation agrees with the previous reports such as [19].After the McColo shutdown, the breakdown of countries changed slightly.

Figure 5 shows the cumulative fraction of hosts in the top countries and ASNs. We notice that the host distributions are more skewed to the top countries/ASNs before the McColo shutdown. Also, the Srizbi hosts were spread over more ASes before the McColo shutdown. That is, before the McColo shutdown, the total number of observed ASNs are 6364 (CORP), 7197 (UW), and 4477 (WAS). After the McColo shutdown, the total number of observed ASNs are 2279 (CORP) and 1290 (WAS). Thus, Srizbi V2 has been less spread, compared to V1 before the McColo shutdown.

3.6 Lessons Learned

From our large-scale empirical study, we derive the following lessons pertaining to thwarting botnet-related spam:

- Attacking C&C servers of the top spamming botnets is an effective way to mitigate a significant amount of spam messages.
- Keeping track of C&C servers is a challenging research task because of the rapid response by the automated fallback mechanisms used by botnets today.
- The contribution of spam from a particular spamming botnet largely differs among receiver domains.

Table 6: Top 10 country codes (CC) of observed Srizbi V1 hosts in Pre-McColo period.

UW Apr 2008		CORP Apr 2008		WAS Oct 2008	
CC	%	CC	%	CC	%
CN	13.3	TR	13.9	RU	21.2
TR	11.8	CN	10.7	TR	16.3
RU	10.7	RU	9.3	IN	6.5
BR	7.2	BR	6.6	US	5.8
US	5.4	US	4.8	BR	4.9
PL	4.1	TH	4.6	KR	3.7
CO	3.7	PL	4.6	UA	3.3
IN	3.7	CO	4.6	CN	3.1
AR	3.3	AR	4.0	PL	2.8
IT	3.3	IN	3.7	TH	2.2
Other	33.5	Other	33.2	Other	30.2

Table 7: Top 10 country codes (CC) of observed Srizbi V2 hosts in Post-McColo period.

CORP Mar 2009		WAS Mar 2009	
CC	%	CC	%
BR	9.5	TR	12.4
TR	7.9	IN	11.0
KR	7.2	SA	7.3
IN	6.7	RU	7.3
CN	5.5	BR	7.0
VN	5.3	CN	4.1
US	4.9	TH	3.3
PL	4.2	PL	2.8
SA	3.8	UA	2.4
RU	3.4	US	2.2
Other	41.6	Other	40.2

- Terminating a spamming botnet permanently requires a long-term and holistic approach.

We believe that these lessons are essential to reducing the spam seen by botnets which have replaced Srizbi. As part of our future work we aim to apply what we know about Srizbi to help gain understanding about the characteristics of these new spam threats.

4. RELATED WORK

Botnets have emerged as a major tool for sending spam from end-host machines. To understand the whole picture of spamming botnets, it is crucial to identify hosts infected with spamming bots. Our work leverages TCP fingerprinting to identify hosts infected with Srizbi botnet without their knowledge for analysis. In this section, we review prior studies that identify spamming bots, and compare them to ours. We then review several studies that leverage TCP fingerprinting to understand the characteristics of spam senders. Along with these studies, our work adds to the growing literature on botnets and sheds new light into botnet mechanics, growth and reaction to counter-measures.

Ways to identify spamming bots have been explored in [5, 16, 18, 20, 23, 25]. Ramachandran et al. [18] developed techniques to identify spamming botnets using passive analysis of DNSBL lookup traffic. The key idea is to find *reconnaissance* lookups from bots. Chiang and Lloyd [5] similarly identified bots, but by monitoring the communication channel between infected hosts and the C&C server of the botnet. Xie et al. [23] developed a framework that

outputs high quality regular expressions that can detect messages coming from botnets accurately. Their method successfully identified 7,721 botnet-based spam campaigns, which utilized 340,050 unique IP addresses from a three-month sample of e-mail messages from Hotmail. Also utilizing a Hotmail data set, Zhuang et al. [25] developed a novel technique to extract botnet membership through the analysis of e-mail message characteristics. By identifying common characteristics, e-mails can be associated with messages of the same spam campaign.

While the characteristics of spamming botnets have been explored in the previous studies, we build upon this knowledge by exploring a particular spamming botnet in detail and analyzing the effect of the take down of its C&C servers. Ramachandran et al. [16] monitored DNS queries to the domain hosting the C&C servers of the spamming botnet, Bobax [8], and discovered around 100,000 bot-infected hosts over 46 days. They studied the completeness and responsiveness of popular DNSBLs using the derived IP addresses of the hosts. Stern [20] studied the architecture of "Reactor Mailer", which is a piece of spamware associated with the Srizbi botnet. Through the careful analysis of a large number of hosts infected with Srizbi, Stern was able to discover the three TCP fingerprinting signatures associated with the botnet. Using the signatures, they successfully connected several significant events involving the spamming botnet. We note that our study reveals several other signature variants of the spamming botnet, which also contribute a large portion of spam messages. Our study also reveals the rapid and drastic transition of the versions of the botnet, triggered by the shutdown.

Finally, we review other studies that leverage TCP fingerprinting techniques to study the properties of e-mail senders [7, 12, 17]. Ramachandran and Feamster [17] analyzed SMTP traffic delivered to their *spam sinkhole* server and found that approximately 95% of the identified spam-sending hosts were running Windows. Similarly, a study by Li et al. [12] investigated the operating system information of the spam host machine, using TCP fingerprinting. They found that 74% of the total spam messages were sent from Windows, around 10% were from Linux, about 5% originated from BSD and Solaris machines, and about 11% were from unclassified hosts. Esquivel et al. [7] et al. proposed a router-level spam filtering scheme that leverages TCP fingerprinting. They showed by targeting the top 100 TCP fingerprint signatures, roughly 60% of spam messages can be eliminated with false positive ratio less than 0.05%.

5. CONCLUSIONS

The temporary but great success of the McColo shutdown indicates the need for a better understanding of spamming botnets as a whole, and the way in which they make transitions is crucial to building effective and sustainable anti-spam solutions. As a first step toward this goal, we studied the world's worst spamming botnet, Srizbi, and the effectiveness of targeting the C&C servers of the botnet, from the viewpoint of several Internet edge sites. We also looked at the long-term trends of the spamming botnet to study how it has grown and changed.

First, we found that the shutdown was actually effective in reducing the volume of spam at Internet edge sites. For CORP and GEM the spam volume seen at these sites was reduced by roughly 40-50% and that reduction lasted at least 2–6+ months. We also found that the long-term effectiveness of McColo shutdown varies at Internet edge sites. We conjecture the difference reflects the way in which the e-mail addresses have been harvested by spammers. Finally, our long-term data set analysis revealed several useful findings in understanding how long a spamming botnet could be active and how fast spammers make transitions between spamming botnets. Our analysis revealed the onset, rise, and steady growth of the Srizbi botnet over the 14 months until the shutdown, and the rapid version transition, triggered by the shutdown.

Our findings suggest targeting a specific set of C&C servers may not be a permanent solution, but it is an effective way to mitigate a significant amount of spam messages at least temporarily. Analyzing the effect of the shutdown is also meaningful to study the way spammers react to action against them. Employing new actions against the infrastructure of spamming botnets, combined with other methodologies, could eventually provide more insights into the tricks used by spammers and narrow their options for recovery. We believe that keeping an ongoing long-term measurement and conducting periodic analyzes is a promising approach for identifying upcoming spamming botnets, studying how they are mitigated by actions taken against them, and building a methodology to stop spamming botnets permanently. Correlating data sets collected at different layers/locations will play a crucial role in understanding the whole picture of spamming botnets.

Acknowledgements

We thank Mike Blodgett, Jesse Thompson, Jeffrey Savoy, and Dave Plonka for their assistance in collecting UW data sets, Mitsuhiro Shigematsu, Masashi Mitsuda, and Kanako Nozue for their assistance in collecting CORP data sets, Wei Yuan, Hiroyuki Ishihara, Yuya Motojima, Junki Ohmura, Kazuhiro Tobe, and Takayuki Tanazawa for their assistance in collecting WASEDA data sets, and Kazunori Takahashi, Noriaki Inoue, Kazuto Noguchi, and Hisao Uose for their assistance in collecting the data set at GEMnet2 [22]. Finally, we thank people of the MAWI WG [13], WIDE project for sharing invaluable data with the research community.

6. REFERENCES

[1] Barracuda Networks Predicts Spam Volumes Beyond 95 Percent in 2009. http://www.barracudanetworks.com/ns/news_and_events/index.php?nid=322, December 2008.

[2] Hosting firm takedown bags 500,000 bots. http://www.computerworld.com/s/article/9120727/Hosting_firm_takedown_bags_500_000_bots, November 2008.

[3] Kaspersky Lab analyses new version of Kido (Conficker). http://www.kaspersky.com/news?id=207575791, April 2009.

[4] Brian Krebs. Host of Internet Spam Groups Is Cut Off. http://www.washingtonpost.com/wp-dyn/content/article/2008/11/12/AR2008111200658_pf.html, November 2008.

[5] K. Chiang and L. Lloyd. A case study of the rustock rootkit and spam bot. In *The First Workshop in Understanding Botnets*, 2007.

[6] Commtouch. Q3 2007 Email Threats Trend Report. http://www.commtouch.com/downloads/Commtouch_2007_Q3_Email_Threats.pdf.

[7] H. Esquivel, T. Mori, and A. Akella. Router-Level Spam Filtering Using TCP Fingerprints: Architecture and Measurement-Based Evaluation. In *CEAS*, 2009.

[8] Joe Stewart. Bobax Trojan Analysis. http://www.secureworks.com/research/threats/bobax/, May 2007.

[9] Joe Stewart. Top Spam Botnets Exposed. http://www.secureworks.com/research/threats/topbotnets, 2008.

[10] Kelly Jackson. Srizbi Botnet Sending Over 60 Billion Spams a Day. http://www.darkreading.com/security/encryption/showArticle.jhtml?articleID=211201479, May 2008.

[11] S. Lawrence and C. L. Giles. Searching the world wide web. *SCIENCE*, 280(5360):98–100, 1998.

[12] F. Li and M.-H. Hseih. An empirical study of clustering behavior of spammers and group-based anti-spam strategies. In *Proc. CEAS 2006: Third Conference on Email and Anti-Spam*, 2006.

[13] MAWI Working Group Traffic Archive. http://mawi.wide.ad.jp/mawi/.

[14] Microsoft. Malicious Software Removal Tool. http://www.microsoft.com/security/malwareremove/default.mspx.

[15] Official Google Enterprise Blog. Spam data and trends: Q1 2009. http://googleenterprise.blogspot.com/2009/03/spam-data-and-trends-q1-2009.html, 2009.

[16] A. Ramachandran, D. Dagon, and N. Feamster. Can DNS-based blacklists keep up with bots? In *Proc. CEAS 2006: Third Conference on Email and Anti-Spam*, 2006.

[17] A. Ramachandran and N. Feamster. Understanding the network-level behavior of spammers. In *Proc. ACM SIGCOMM 2006*, pages 291–302, 2006.

[18] A. Ramachandran, N. Feamster, and D. Dagon. Revealing botnet membership using DNSBL counter-intelligence. In *2nd Steps to Reducing Unwanted Traffic on the Internet Workshop (SRUTI)*, 2006.

[19] Sophos. Security threat report:2009, Jan 2009.

[20] H. Stern. The rise and fall of reactor mailer. In *Proc. MIT Spam Conference 2009*, Mar 2009.

[21] The Spamhaus Project. The Policy Block List. http://www.spamhaus.org/pbl/index.lasso.

[22] H. Uose. GEMnet2: NTT's New Network Testbed for Global R&D. In *TRIDENTCOM '05: Proceedings of the First International Conference on Testbeds and Research Infrastructures for the DEvelopment of NeTworks and COMmunities*, pages 232–241, 2005.

[23] Y. Xie, F. Yu, K. Achan, R. Panigrahy, G. Hulten, and I. Osipkov. Spamming botnets: signatures and characteristics. In *Proc. ACM SIGCOMM 2008*, pages 171–182, 2008.

[24] M. Zalewski. the new p0f: 2.0.8. http://lcamtuf.coredump.cx/p0f.shtml, 2006.

[25] L. Zhuang, J. Dunagan, D. R. Simon, H. J. Wang, and J. D. Tygar. Characterizing botnets from email spam records. In *Proc. USENIX LEET 2008*, pages 1–9, 2008.