Phishing E-Mails – Six Month Investigation into What People Click

Michael R. Lehrfeld
Department of Computing
East Tennessee State University
Johnson City, TN 37614
423-539-6952
Lehrfeld@etsu.edu

Abstract

Phishing and SPAM emails have been used by marketers and hackers alike since the inception of email and the Internet. Phishing messages have become so common that many legitimate emails often get flagged and placed in a user's spam bucket. No one is denying that these messages are at a minimum a nuisance, and in many cases malicious. But what is known about the success rates of these phishing campaigns? What are users most likely to fall victim too? Why do victims fall prey to these types of messages? This research attempts to shed some light on these questions.

During a six month period, over 33,500 phishing emails were sent. Each of these emails was tagged with characteristics from a developed phishing matrix do determine which phishing messages are more likely to be clicked. The phishing matrix categorized messages into 4 dimensions; time, deception, geography, and group affiliation. What was discovered is that the most successful dimensions were time and deception. Phishing emails that exploited current events received high click rates as well as those with colorful figures and backgrounds. It was found that 15% of users clicked on more than one phishing campaign leading to an addicted clicker syndrome; a situation where a user will click on almost any message that enters their inbox.

Introduction

Some of the most prolific and damaging data breaches have been attributed to targeted phishing campaigns against employees and companies[1] [2]. The crown jewels of technology companies have been pilfered from the most secure systems; and in the case of the RSA data breach, the very companies whose business is the protection of secure systems. Despite the massive outlay of capital to protect computing systems hackers are still able to find weaknesses in the defenses and gain access to sensitive information.

The root at some of these high profile breaches has been phishing emails. Phishing emails are defined by the unsolicited receipt of emails that attempt to coerce the victim into purchasing a good or service, clicking on a link that directs a user to a malware serving website, download and run a harmful attachment, or have the victim respond to the email with their user credentials or banking information. In all cases, the victim must act for the phishing campaign to be successful. The simple receipt of the message is not enough to cause harm.

This research seeks to discover the most effective types of phishing emails so that targeted training can be developed to help prevent users from falling victim. To facilitate this goal, phishing emails have been divided into four main dimensions; each of which contain three scoring levels as to the phishing email sophistication within that particular dimension. Over a six month period, 33,500 phishing emails were sent to employees of a large company. Twelve phishing campaigns were successfully delivered while four were rejected because of technical issues with the organizations spam filters. Over the testing period, 1396 phishing emails were clicked by the users.

To protect the identity of our population, all email addresses and identifying information has been obfuscated. The remainder of the paper is composed of a summary of the current state of phishing detection and prevention techniques, the methodology that was used in the development of the phishing dimensions and the Phishing Matrix, an overview of the phishing architecture used to send the phishing emails, a discussion of the results of the 12 phishing campaigns, and finally a conclusions and future works section.

Literature Review

Phishing and SPAM emails have been a by-product of the Internet age since its inception. Unsolicited emails selling discount medical cures, promises of wealth, and scare tactics have been some of the strategies that spammers employ to get unsuspecting users to click on their emails. The proliferation of email in both the business and personal environment as a primary vehicle for communication just enlarges the victim pool for spammers. Some accounts place the global internet email address pool at over 3 billion with an anticipated additional 1 billion new accounts by 2017 [3].

The global email numbers continue with staggering volume. On average, over 120 billion email messages are sent daily [4]. Radicati [3] estimates that each email account is responsible for over 100 emails per day. This amount of communication demonstrates the reliance of the medium in modern life. It also represents an easy medium to exploit by hackers.

Many attempts have been made to prevent phishing emails from ever reaching the inbox of users. The spam botnet research [5] examined the characteristics of over 50 million spam messages to ascertain campaign duration and depth. This research discovered a sophisticated relationship between phishing campaigns and the botnet workloads. This complex interworking leads to more intricate spam blocking infrastructure to protect the end users. The work of Chang et al. continues the investigation to defending against web based attacks; including phishing [6]. In their work, the researchers utilize honeypot techniques and web site blacklisting and analyze their effectiveness in disrupting internet attacks. What can be drawn from these studies, and others like it [7], is that blocking 100% of phishing emails is currently an impossible task.

The following section discusses the development of the phishing matrix that is used to categorize phishing campaigns and their content.

Methodology

Assuming that some phishing emails make it to their destination, this research sought to examine what, if any, characteristics made particular phishing emails more or less likely to be clicked on. Relying on the user base to give an accurate account of why they have fallen victim to phishing emails gave some insight into the problem. A 2010 study by the Messaging Anti-Abuse Working Group (MAAWG) [8] surveyed 3700 people about how they determined if an email was legitimate or a phishing attempt. The top indicators were 1) sender’s address (73%), 2) subject (67%), unusual language (53%), content of the message (53%), salutation (46%), spelling and grammar (43%), and visual cues in the message (30%). The MAAWG study also discovered that over 11% of the respondents had clicked on phishing emails in the past. It was these self-reported triggers that the phishing matrix was built from and elaborated upon.

Using the MAAWG study as a starting point, the phishing matrix was developed. It included four general dimensions (or categories) of phishing emails. Each dimension contains three attributes that measure the severity of the applied dimension. The reminder of this section discusses the four phishing dimensions and their associative attributes.

Group Affiliation Dimension

The first phishing dimension that is examined is group affiliation. This dimension targets a victim's identity with a particular sub-group. The purpose of this measure is to determine if targeted phishing campaigns that associate with common groups need further investigation. Group affiliation content is focused on broad, well established populations. Types of group affiliation phishing campaigns would be "all retirees", "all parents", or "healthcare recipients".

Within this dimension, there are three attribute levels employed; low, medium, and high. In order for a message to be classified as 'high' the content of the message must explicitly reference a group type (e.g. all .EDU users, all Medicare recipients). Medium classification implies an association with the target group while low removes all group markers. This classification may limit the potential users that would click on the spam but yield a higher click rate within the classification.

Deception Dimension

Phishing, by its nature, is a deceptive practice. However, this dimension is focused strictly on the technical sophistication of the phishing campaign. The majority of email users self-identify with being 'somewhat', 'very experienced', or 'expert' when asked about their knowledge of Internet security [8]. This dimension seeks to challenge those findings by using simple HTML tables in phishing messages.

The deception dimension was divided into three attributes of low, medium, and high. Low represented no deception at all. The *From* address was not obfuscated and no graphics or coloring were used. The medium attribute incorporated some color or figures and minor changes to the URL were used but remained primarily unchanged from the low attributed message. The high attribute attempted to trick the user by presenting the *From* address as legitimate, using more complex figures and coloring, or altering the URL in a more complex manner. The salutation and closing of the message may also have been more familiar to the user rather than a generic one.

Geographic Dimension

Classifying phishing campaigns with the geographic dimension utilized the victim's affinity with a geographic area. This dimension examined if a user received a phishing email that referenced a location that they were familiar with, would they be more likely to fall for it? This research utilized the fact that the target population of this study was mostly isolated to 30 mile radius.

A three point scale is again used with the low, medium, and high attributes. The low attribute refers to the absence of any mention of a geographic element in the message. The next level of the geographic attribute is medium. This attribute is defined by a larger area than the high attribute and encompasses language such as USA, England or France. The high attribute refers to a highly recognizable geographic location by the recipient of the phishing message. This type of message contains language such as "the best car dealership in the North East" or "everyone living in East Tennessee should be aware of this". The geographic attribute is bound by regional elements as the most specific area in this study. It was determined that using a city was too close to spear phishing and thus not included.

Time Dimension

The Time dimension examined click rates of messages that incorporated current events. Time is the primary motivator for action in this phishing messages. This dimension is categorized by either 'countdown' events like "you have 3 days until your email is deleted" or current events such as "Payton Manning set to retire" being sent just after the 2014 Super Bowl.
Attributes of the Time dimension are in the three point scale of low, medium, and high. The low attribute is applied when the phishing email is devoid of any reference to time. An example of this is a message asking the user to click on a link to see pictures of cute kittens. The medium level time attribute includes an indirect reference to time as a motivator. Time is not explicitly stated but the context of the message relays a time component. An example of this type of spam would be a spam message about textbook leasing at the beginning of a term. The high attribute occurs when there is an overt and direct reference to time as a motivator for clicking. Examples of these include countdown content like password resets within 7 days, or messages about Christmas or other holidays.

Phishing Architecture

This section discusses the anatomy of the emails that are sent to the targets, how the emails are tracked, and how reporting is generated using the email architecture established in this section. The emails are constructed in a way to ensure the most effective reporting and tracking possible.

The email architecture consists of a combination of XML and HTML. The XML is used to create embedded tags that allow each campaign to be tracked individually. Each email is given either a low, medium, or high XML tag to help distinguish it from the other urgency levels within the phishing matrix. The emails are individually named as well using a convention like: fc11g01. This denotes which campaign it is and which urgency level as well. For example, campaign 11 would consist of 3 different emails with the following naming convention: fc11g00, fc11g01, fc11g02. The incremental value, within the name, denotes the urgency level or complexity of the email (fc11g01, would be campaign 11 with urgency level medium). The campaign attributes are also included within the email to track emails based on which dimension they fit in. See section 3 for a discussion of the phishing dimensions.

To protect the identity of the phishing email recipients, all email addresses are uniquely hashed prior to the message being sent. This allows for the insertion of a unique ID (GUID) for every user, for every campaign, in the emails. The format for the link phishing victims' click takes the form in Figure 1. Using the unique GUID ensures that if the emails are coopted or the web server is compromised, no identifiable information can be obtained.

http://retirement.albybum.org/fc21g01/?b580015f20674154a09b81b063438460

Unique GUID Campaign Identifier URL

Figure 1. Phishing link construct.

The HTML portion of the email is used to house the content that is being sent to a target. This includes: subject, content, inline CSS, and various HTML tags. When the email is created, depending on the level of urgency, subject lines can be modified for each email within each campaign. This allows for the raising and lowering of suspicion within the email itself. Content and visual styling can also be modified for each email within a campaign to make it more attractive to targets.

For each campaign, 3 phishing emails are created that are divided up into low, medium, and high urgency levels. This allows for 3 modified versions of the same email to track different responses based on style or wording within the email itself. As the email increases in urgency the complexity increases as well. The victims are divided up into 3 equal bin. This allows for two important measures; 1) specific modifications to a base phishing email to test only one dimension and, 2) decreases the probability that people in close physical proximity don't receive the exact same phish. For each campaign the emails are shuffled amongst the targets to ensure that the same group did not get the same level each time.

When an email is sent, the content of the email, including the XML file itself, is stored in a database. This database maintains a record of when the emails were sent, and the content within each email. This database resides on the web server and does not contain any identifying information about the recipients of the phish.

The reporting of the data captured from the spam clickers is broken up into a csv file using a Python script. Once the data has been captured, the script is run against the captured database. Using the XML tags mentioned above a csv file is produced that breaks each click into components. An example of different components would be time of day, browser type used, and the user's email address. At the moment when the report is created using the Python script, the target's GUID merges with the target's actual email address. Prior to this engagement, the target's name and email address are kept separate from each other by residing on two separate servers.

The implementation and testing of the spam matrix falls into three subsystems. These are discussed in further detail in the following subsections; 1) storage and tagging phishing messages, 2) sending and tracking phishing messages, and 3) reporting.

Phish Message Categorization Recording Method

The phishing message creation processes begins by storing the messages in an XML schema that contains the tags and overall structure of the email. This XML document encompasses HTML and inline CSS for styling. For each XML document there exists only one email, so one file equals one email. This allows for multiple messages to be developed in parallel as well keeping a history log of every email that is sent out. This architecture reduces the overall amount of code required to interact with the email, since each email is similar. Only the email content and tags need modification.

To document the schema and to gain better recording of the data, the XML schema is broken down into metadata tags. The first metadata tag is the Title; this simply describes the nature of the email so it can be differentiated from the other emails. The second metadata tag is the Author; this exists to track the person responsible for creating the actual email. The third metadata tag is the Date; this describes the date in which the email was created by the previously mention author. Last is the Matrix tags; this encompasses various tags that are used to actually track the clicks in a campaign.

The next schema that is defined is the email data schema. The email data schema defines the actual email itself and is broken down into its appropriate categories. The first category is the Subject; this is simply the subject of the email. The second category is the Body; this is where the message content resides. The Body tag also holds the links in which the target would click. The third category is the From Address; this is the address in which we want the email to appear to come from. The From Address would ideally relate to the content of the body. Last is the Fake SMTP server; this is the server in which we use to show where the actual email came from. Once again, this can be anything but ideally it should relate to the topic of the email.

The XML document modification process can be quite complex at times. Therefore, certain fail safes have been put into place to accommodate for human error. In terms of the XML document, a separate Python script has been created to validate the layout to ensure that it meets the predefined schema. If the schema is not validated, the sending script is stopped, and the email is not sent to ensure data integrity.

Sending Phishes

The emails are sent via a python script that queries the target database. It sends the emails to the appropriate target groups (group 0, group 1, group 2), defined by the user, when the user executes the python script. There is a collection of Python scripts responsible for taking an XML spam message and sending it to a group of targets. There are three main scripts:

33. Script for creating the admin database. All data is stored in a SQLite database. A database represents a collection of spam campaigns. This Python script creates all database tables needed for sending spam campaigns.
34. Script for loading targets into the database, as well as grouping them for campaigns. The number of groups is specified when the database is created. New targets can be added at any point in time—this will not break reporting. Additional information may be included about targets for reporting purposes.
35. Script for sending the emails. This script takes an XML phishing message and sends it to a specified group. The email’s subject, from field, and body are templates, and may contain placeholders that pull information from the target’s information. For each email sent, a GUID is generated to uniquely couple the target and email sent.

Tracking Phishes

The phishing email is built to be tracked using the XML schema described above. Within this XML document, XML tags are used to allow for the campaign to be linked back to the target who clicked it. The XML tags are also used to track the amount of clicks, the campaign, the content sent, metadata, and the GUID. The GUID provides a way to link the email back to the target who clicked the link.

The modular approach to the phishing system has allowed the transmission of some sophisticated emails. As privacy is a top concern in this research, the ability to physically separate the user's email addresses from the web server is paramount to this project moving forward. Even if the public facing web server is compromised, the identities of the recipients is still not in jeopardy. The following section will discuss the results of the phishing campaigns followed by a conclusions section.

Results

Over 33,500 phishing emails were sent during a six month period that saw 1,396 of those messages clicked. The 33,500 phishing emails were broken into 12 different campaigns across the phishing dimensions. Table 1 and Figures 2 and 3 presents the general breakdown of clickers by campaign.

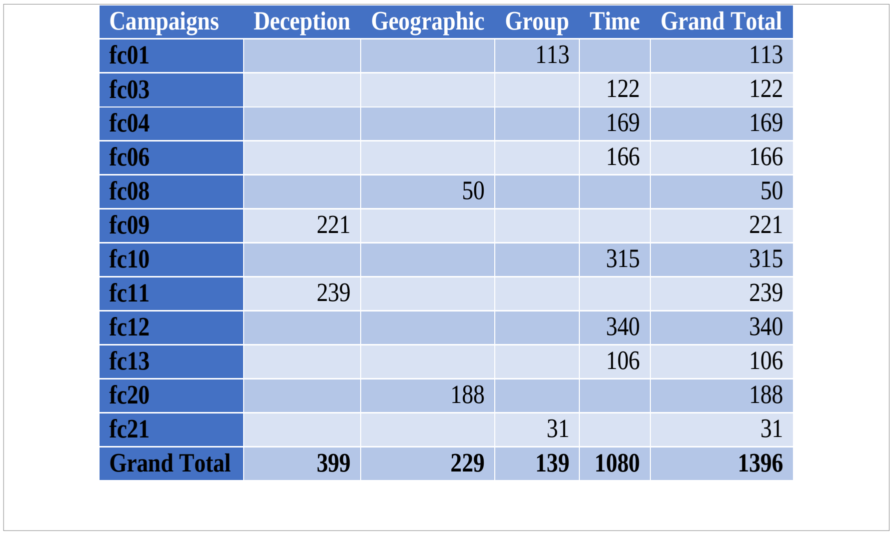

Campaigns	Deception	Geographic	Group	Time	Grand Total
fc01			113		113
fc03				122	122
fc04				169	169
fc06				166	166
fc08		50			50
fc09	221				221
fc10				315	315
fc11	239				239
fc12				340	340
fc13				106	106
fc20		188			188
fc21			31		31
Grand Total	**399**	**229**	**139**	**1080**	**1396**

Table 1. General statistics about phishing campaigns and click rates.

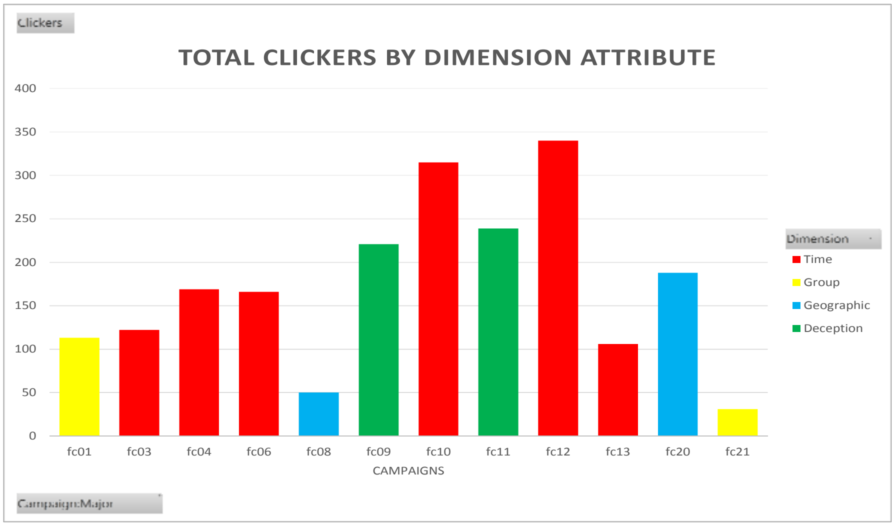

Figure 2. General statistics about phishing campaigns and click rates.

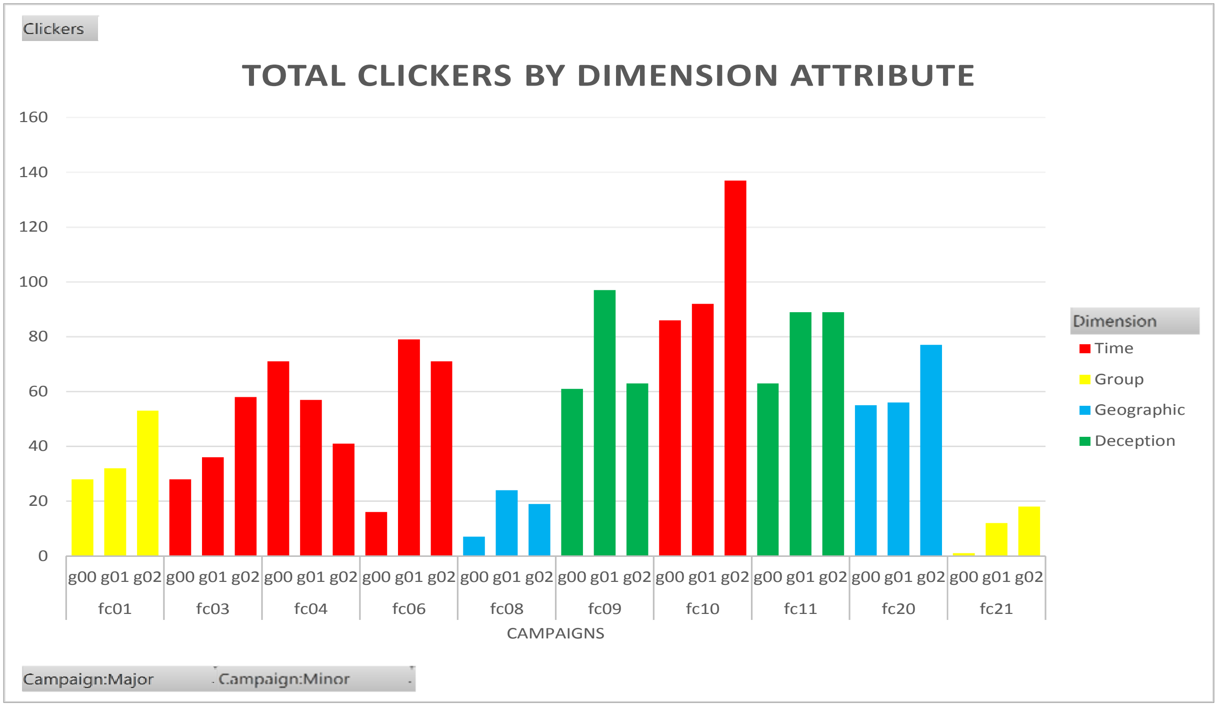

Figure 3. Click rates by dimension attribute. Note, campaigns fc12 and fc13 did not have capture attribute information.

An evaluation of people that clicked on multiple campaigns was also performed. This group was dubbed ‘Addicted Clickers’ and is represented in Figure 4.

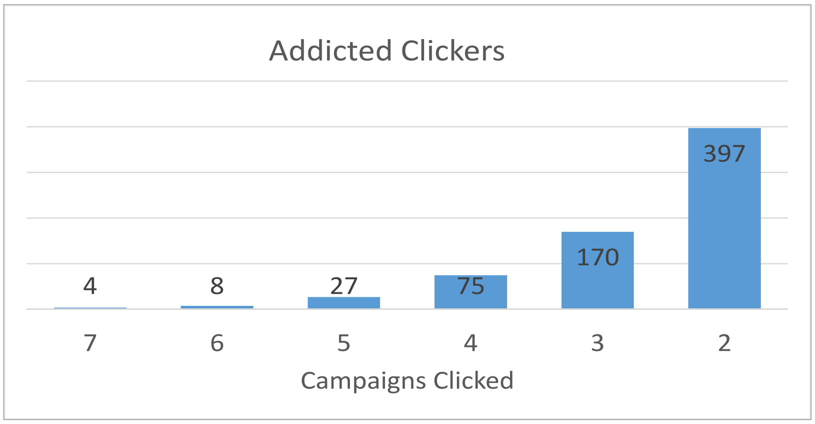

Figure 4. Number of people that clicked on multiple campaigns.

The following section discusses conclusions drawn from the phishing matrix data, some lessons learned from the project, and future works.

Conclusions and Future Work

People love to click on phishing emails. At the inception of the project, a baseline email was sent out that just said ‘Click Here’. 183 people did. 15% of the target population clicked on two or more phishing emails. The most effective dimension was Time followed closely by Deception. Of all the campaigns, the one campaign that generated the most clicks was a phishing email sent at the end of a month stating that a system upgrade has caused issues with their direct deposit. Not only did this set the record for clicks, but as our landing page for our clickers was a 404 error, many concerned people called the payroll department in a panic because the website was ‘down’.

As seen in Figure 3, many of the campaigns did not show much of a clicker variation between the medium and high attribute messages, and in campaign 4 the low message had the highest click rate. More phishing needs to be conducted within each of the 4 dimensions to get a better understanding of why this is.

Future areas that will be evaluated include the addition of a Greed dimension. This dimension will accompany the current 4 other dimensions and focus on campaigns around winning a prize, super sales occurring, or others attributing factors where the victim perceives some type of financial reward from clicking. Another area that will be researched is the enrichment of the current dataset with user demographic data. The ability to categorize clickers by age, gender, race, employment position, and salary will give greater insight into the data and what groups are clicking the phishing messages. Finally, starting in the fall, the project will expand to include 65,000 student email accounts.

The phishing architecture will be overhauled to enable a larger set of phishing emails to be sent. Data on campaigns that ask the users for their userID and password need to be investigated. The system will also be enhanced to monitor phishing email that has been forwarded to others. For example, if there is a phish about breaking news, do people forward those messages to others.

Lastly, all of this data collected will be used to feed the development of a training program. This program will seek to keep people safe at home and at work and will target the areas of the phishing matrix that have produced the most clicks by users.

References

[1] J. Reid and K. Coratti, "Washington Post Site Hacked After Successful Phishing Campaign," vol. 2014, ed, 2013.

[2] RSA FraudAction Research Labs, "Anatomy of an Attack," in *RSA Speaking of Security* vol. 2014, ed, 2011.

[3] Radicati, "Email Statistics Report, 2013-2017," Radicati Group2013.

[4] C. Kanich, C. Kreibich, K. Levchenko, B. Enright, G. M. Voelker, V. Paxson, *et al.*, "Spamalytics: an empirical analysis of spam marketing conversion," presented at the Proceedings of the 15th ACM conference on Computer and communications security, Alexandria, Virginia, USA, 2008.

[5] A. Pathak, F. Qian, Y. C. Hu, Z. M. Mao, and S. Ranjan, "Botnet spam campaigns can be long lasting: evidence, implications, and analysis," presented at the Proceedings of the eleventh international joint conference on Measurement and modeling of computer systems, Seattle, WA, USA, 2009.

[6] J. Chang, K. K. Venkatasubramanian, A. G. West, and I. Lee, "Analyzing and defending against web-based malware," *ACM Computing Surveys,* vol. 45, pp. 1-35, 2013.

[7] T. A. d. O. Alves and H. T. Marques-Neto, "Characterizing a Network of SPAMs Recipients," presented at the Proceedings of the 2012 Eighth Latin American Web Congress, 2012.

[8] Messaging Anti-Abuse Working Group, "2010 MAAWG Email Security Awareness and Usage Report," ed, 2010.