An analysis of the effectiveness of personalized spam using online social network public information

Enaitz Ezpeleta[1], Urko Zurutuza[1], and José María Gómez Hidalgo[2]

[1] Electronics and Computing Department, Mondragon University
Goiru Kalea, 2, 20500 Arrasate-Mondragón, Spain
{eezpeleta,uzurutuza}@mondragon.edu,
[2] Pragsis Technologies
Manuel Tovar, 43-53, Fuencarral - 28034 Madrid, Spain
jmgomez@pragsis.com

Abstract. Unsolicited email campaigns remain as one of the biggest threats affecting millions of users per day. Spam filters are capable of detecting and avoiding an increasing number of messages, but researchers have quantified a response rate of a 0.006% [1], still significant to turn a considerable profit. While research directions are addressing topics such as better spam filters, or spam detection inside online social networks, in this paper we demonstrate that a classic spam model using online social network information can harvest a 7.62% of click-through rate. We collect email addresses from the Internet, complete email owner information using their public social network profile data, and analyzed response of personalized spam sent to users according to their profile. Finally we demonstrate the effectiveness of these profile-based templates to circumvent spam detection.

Keywords: spam, security, Facebook, personalized spam, online social networks

1 Introduction

The mass mailing of unsolicited e-mails have been one of the biggest threats for years. Spam campaigns have been used both for the sale of products such as online fraud. Researchers are investigating many approaches that try to minimize this type of malicious activity that report billionary benefits.

Within the spam problem, most research and products focus on improving spam classification and filtering. According to Kaspersky Lab data, the average of spam in email traffic for the year 2014 stood at 66.9% [2].

With the rise of online social networks (OSNs), specifically Facebook, which has more than 1.39 billion monthly active users as of December 2014 [3], the extraction of personal information that users leave public on their profiles multiplies spam success possibilities. Facebook provides a great opportunity for attackers to personalize the spam, so a much lower volume of messages would get a higher return on investment.

The main objective of this paper is to measure the consequences of displaying information publicly in OSNs. It also aims to demonstrate that advanced techniques for generating personalized email that evade current spam detection systems while increasing the click-through rate can be developed. These techniques can enable new forms of attacks. First we extracted email addresses while crawling the Internet. These addresses were then checked on Facebook to look for related profiles. Once stored a considerable quantity of user addresses, we extracted all the related public profile information and temporally stored it in a database. Then this information was analyzed in order to determine user profiles based on their main activities in Facebook. Email templates were generated using common information patterns. Finally, to demonstrate the effectiveness of these templates when systems circumvent spam detection, different experiments have been performed. We collected sufficient evidence to confirm that the goal has been achieved.

The remainder is organized as follows. Section 2 describes the previous work conducted in the areas of personalized spam, and social network spam. Section 3 describes the process of the aforementioned experiments, regarding data collection, data processing, and personalized spam testing. In Section 4, the obtained results are described, comparing typical spam results with the personalized ones. Section 5 gives a discussion of the countermeasures that can be applied to prevent personalized spam. Section 6 describes the ethical considerations about this research. Finally, we summarize our findings and give conclusions in Section 7.

2 Related Work

2.1 Personalized spam

During the last years several works about the possibilities to create personalized spam or collect personal information from different OSNs have been proposed.

In 2009, researchers at University of Cambridge and Microsoft analyzed the difficulty of extracting user information from Facebook to create user profiles [4]. They described different ways of collecting user related data, and they demonstrated the efficiency of the proposed methods. Authors conclude that the protection of Facebook against information crawlers was low. They also they proved that big scale collection of data is possible. While it is true that Facebook has improved its systems' security since then, like limiting its own query language, the research proved that the option was effective.

In [5] researchers found a Facebook vulnerability giving attackers the possibility of searching people through email addresses in OSNs. Starting from a list of different emails, they managed to connect those email addresses with the account of their owners. After that, they collected all the information they could, and created different user profiles. This work left open doors for allowing attackers to launch sophisticated and specific attacks, but still did not realize about the potential of creating personalized spam campaigns. In the same direction, Polakis et al. demonstrated in [6] the risk that different OSNs suffer to create personalized spam campaigns.

2.2 Online social network spam

Over the last few years, social networking sites have become one of the main ways to keep track and communicate with people. Sites such as Facebook and Twitter are continuously among the top 10 most-viewed web sites on Internet [7]. The tremendous increase in popularity of OSNs allows them to collect a huge amount of personal information about users. Unfortunately, this wealth of information, as well as the ease with which one can reach many users, also attracted the interest of malicious parties.

Researchers from the University of California proved that spam is a very big issue for OSNs [8]. In their research they created a large and diverse set of false profiles on three large social networking sites (Facebook, Twitter and MySpace), and logged the kind of contacts and messages they received. They then analyzed the collected data and identified anomalous behaviors of users who contacted their profiles. Based on the analysis of this behavior, they developed techniques to detect spammers inside OSNs, and they aggregated their messages in larger spam campaigns. Results show that it is possible to automatically identify accounts used by spammers, and block these spam profiles.

In Gao et al. [9] authors carried out a study to quantify and characterize spam campaigns launched from accounts on OSNs. They studied a large anonymized dataset of asynchronous "wall" messages between Facebook users. They analyzed all wall messages received by roughly 3.5 million Facebook users, and used a set of automated techniques to detect and characterize coordinated spam campaigns. This study was the first to quantify the extent of malicious content and compromised accounts in a large OSN. While they cannot determine how effective these posts are at soliciting user visits and spreading malware, their result clearly showed that OSNs are now a major delivery platform targeted for spam and malware. In addition, their work demonstrates that automated detection techniques and heuristics can be successfully used to detect social spam.

While most of the research focus on spam campaigns that might appear inside OSNs, we still think that a combination of typical spam and OSN spam exposes serious threats that needs to be addressed.

3 Creating a personalized spam campaign

Our study was carried out in four different phases. In the first phase, we collected a large amount of public information from Facebook. To do this we used email addresses that were publicly available when crawling the Internet. In a second step of our research, we computed a number of interesting statistics from the collected information that will be shown later. As a result of the data analysis, different user profiles were identified, and used then as customizable email template. Once we had defined these templates, we developed an automatic email sending system and conducted two different experiments. Finally we analysed the results obtained in the experiments.

3.1 Collection of data

At this stage, the main objective has been the collection of information. This process have been performed in three steps:

Email address collection. In this task we considered two options: the first one, obtain the email addresses using the techniques as explained in [6], where they get e-mail addresses using various combinations of public information from users OSNs' users. And the second, using applications to find emails on the Internet. In the end, we have chosen the second one.

Email address validation. First of all we have worked to obtain the greatest numbers of emails. After that, we have checked each email in Facebook, and we have obtained the amount of emails related with a Facebook account.

In order to carry out this work, we have seen that Facebook offers the option to find a personal account through the email. An application to automate the validation has been developed. It first authenticates a user to the OSN, and then searchers for a user corresponding to each email harvested before.

Once we found the account we have extracted and saved the user's ID and the full name.

Next URL is used to check if an specific email (shown in the URL as EMAIL) corresponds to a Facebook user.

```
http://www.facebook.com/search.php?init=s%3Aemail&q=EMAIL&type=users
```

Collection of the information. On Facebook there is the possibility to extract information from the source code of all the pages. But to do this, it is obligatory to access directly to the page from which we want to extract the information. That is, when you enter Facebook, if you visit the Facebook page of a friend directly, you could get all the user's public information from the source code. On the contrary, if from that page of a friend you visit the profile of a third person, from the source code you cannot get any information. Therefore, in this program we have used user identifier from Facebook to connect directly to the user information page. Thus, we have visited all users pages and we have been able to extract all the public information that users have in the Facebook database.

Below is the address where can be found all public information of each user. 'USERUID' correspond to the ID that Facebook gives each user, which is stored in the database, because we have collected in the previous phase.

```
http://www.facebook.com/profile.php?id=USERUID&v=info
```

Results. We found that a 19% of the email addresses in Internet have a corresponding Facebook account associated to it. We found 22,654 Facebook accounts using 119,012 email addresses (19.04%).

3.2 Data Processing

At this stage the aim has been to treat the data stored in the database to extract user profiles. To facilitate this task we have decided to create a new table in the database. In this table we saved temporally the summarized information for each user. Using the statistical table, we have queried it to extract as much knowledge as possible. After that, we have analysed all the data to present the most important or relevant statistics.

Results. With the previous statistics, we can draw the following conclusions:

- Most Facebook users choose their favourite band and leave it public.
- 30% of users who have some data entered and posted on Facebook, have at least one company with which they have been connected.
- The number of men is 12 percentage points larger than women on our study.

3.3 Personalized spam

The objective of this phase is to create different templates that will be later used. With this templates it is possible to send personalized mail to all Facebook users stored in the database. Once the templates were designed and implemented, the next step was to create a way to count the number of users that "bite the hook" of the spam. For this we have implemented a website.

Mail templates. Before any other action, the first step was to define a template through which we were able to send personalized emails to the people.

To do so, we have taken the decision to create the templates after analyzing the data and the information stored in the database, and using statistics.

For the analysis we used the table in which it was possible to see the number of users that had inserted data in each variable. The table 1 shows the five more introduced variables.

Table 1. Number of users who have entered each variable (total users: 22,654)

Variable	Amount	Percentage
man	8,786	39%
woman	6,189	27%
music	5,788	26%
titles	5,612	25%
company	5,149	23%

As it can be seen, the most abundant variables are those related to the gender. Although these data cannot be used for creating templates, it can be used for implementing a formal greeting according to their gender.

In contrast, we have used other three variables to create spam templates. That is, if a user has entered his favourite music group in the profile, it will receive a personalized **music template**. However, if the person has no any singer or group in Facebook and has added the university in which he or she has studied, it will received a personalized message with the **studies template**. And if none of those two had been added but the information refers to their current job or a company in which the user works, it will received a personalized the email using the **company template**.

For better customization, we have also used some profile fields such as the language, the name of each user, the gender, and the city in some of the templates.

It should also be added that all emails will include a URL through which users access to our website.

Website. Access to this site should only come from the users personal email. It has been necessary to define a system for it. We must also take into account that it should store information about which user, and by what type of mail has come to the site. Considering all these details, we decided that the most appropriate way was to introduce parameters in the URL which will be included in emails. When the user clicks on the URL and gets to the site, these parameters will be stored in the database. It also gives the user the possibility to write a comment or to unsubscribe from the system so that no longer will receive emails.

4 Experimental Results

We have performed two separate experiments. In the beginning, we sent typical spam from a classical spam text in order to measure the success rate. There could be different possibilities here. The spam could have been detected and filtered by the email service, Internet Service Provider, email client in the users computer, or it could have been ignored or deleted by the user. In the second experiment, we focused on personalized spam, in order to prove the click through rate obtained, sending a bigger amount of personalized spam. The results of each experiment, and explanation thereof will be explained in two different subsections. And the comparison of the results in the last subsection.

4.1 First experiment: typical spam

Using multiple email accounts and sending a total of less than one hundred emails per day, we sent a typical spam email. The account change is due to a strategy to make things more difficult to spam detection systems. We sent one of those emails where spammers try to draw the receiver's attention to enter a web address. In total, we sent 972 typical spam emails. Only four users followed our website address. This means that the click-through of the typical spam in our experiment is 0.41%.

4.2 Second experiment: personalized spam

Once extracted the data shown in the previous experiment, it has conducted the second experiment. In this case, instead of sending typical spam, it was sent a personalized email to 2,889 Facebook users. We used the same strategy as in the first experiment sending the messages, and we sent each template from different email accounts. Always taking care to sent less than one hundred emails per day and account.

As we mentioned previously, we use three different templates in our study. Those templates had a personalized URL that could keep track and detail of each sent mail. Note that the website explained the experiment, apologizing for the damage caused, and left space for user comments.

The next table shows the amount of emails we sent, and the number of emails for each profile templates. There we can see that the most common template we used was the once we call 'Music'. More than the 60% of the personalized mails were about music preferences of users. This is mainly because it is the first templates that our program try to use. And if it is not possible to use this template, the application try with the next one (Studies). And the last option is the 'Company' template.

Table 2. Number of sent emails

Type	Amount	Percentage
Music	1,787	61.85%
Studies	843	29.18%
Company	259	8.97%
Total	2,889	100%

To analyze the responses, we must analyze the results automatically stored in the database. As we can see in the table 3, 220 users have acceded the website. This is 7.62% of the people that received a personalized email. Also note that 1.38 percent of people have discharged from the study.

Table 3. Website data

	Amount	Percentage of total shipments
Users who have accessed the website:	220	7.62%
Users who have been discharged:	40	1.38%
Users who have left comments:	11	0.38%

Moreover, we can also break down the answers taking into account the different templates. Because each user acceded to the website from the personalized template as it is shown in the next table together with the click-through of each template.

Table 4. Information according to each template

	Access to the website	Percentage of total accesses	Click-through
Music	111	50.45%	6.21%
Studies	81	36.82%	9.61%
Company	28	12.73%	10.81%

As we can see in the table, most of the users, who entered in our website, had received musical spam. This can be considered as normal, because we sent more musical spam than the other two types. But it is important to see that the template with the highest click-through rate is the 'Company' template. Otherwise, the musical template had the worst results.

4.3 Comparison between experiments

Finally, if we want to analyze the results of the second experiment in the best way, we have to compare with the results of the first experiment. In this way we can see the difference between typical spam and personalized spam results.

Table 5. Comparison between results

	Sent	Answered	Percentage
Typical spam:	972	4	0.41%
Personalized spam:	2,889	220	7.62%

Table 5 summarizes the response rates obtained using the different spam types. If we analyze these data further back, the first interesting information that emerges is that only 4 people have gone through the typical spam. In contrast, 220 other people have come through personalized email. I.e. 0.41 percent compared to 7.62 percent.

5 Countermeasures

Completed the experiments, here are three ways to avoid spam customization. Two from the OSNs point of view, and the other from the users perspective.

- *Limiting users public information:* OSNs may limit public information from users. Thus, will be more difficult to extract information from users. And the attackers can not use this information in their attacks.
- *Changing the code of the website:* Today it is possible to collect information from the source code of the Facebook web page. If they change the website and do not leave the user information in a readable format, it will be more difficult to extract information for attackers.
- *Raising Awareness:* We must teach people how dangerous it can be to leave personal information publicly. If people minimize their profiles public information will be much more difficult to customize the spam.

6 Ethical Considerations

Some actions taken in this paper are ethically sensitive. For some people, collecting information from internet is not ethically correct. But as was said in [10, 11] and more recently in [5], the best way to do an experiment, it is to do as realistically as possible. We defend this mode of action for the following reasons.

First, we must be clear that we work to improve the safety of users, we use users information to protect them in the future. Second, we only use information that users displayed publicly in OSNs. This means that we never attacked any account, password or private area. Third, attackers use this information, if we use the same information and act in the same way, we will defend users better.

Finally, we have consulted to the general direction of our university and they have given us the approval. For this, before the experiment we proposed our intentions to the general direction of the university, where we showed them the ethical considerations for conducting the study. We also explained them the procedure we had designed to collect personal data and the way we had thought to send emails. Once the R&D Manager had gave us the approval, we started with the experiment.

7 Conclusions

This work makes clear the issue that could exist if spam campaign creators turn their spam templates into a personalized text based on user characteristics, interests, and motivating subjects. Attackers have millions of email addresses stored. We have demonstrated that a 19% of the email addresses in Internet have a corresponding Facebook account associated to it. Moreover, even if not too much, basic public information can be extracted from those users, which is sufficient to create personalized email subject and bodies. This emails can have a click-through rate bigger than an 7.62%, being it more than 1,000 times bigger than typical spam campaign rate. It is obvious that in parallel to the research of new techniques for spam detection inside OSNs, it is necessary to research beyond the state of art of classic spam filtering, taking into account the possibility of personalized spam campaign success.

Regarding the behavior of OSN users analyzed, we found that most of Facebook users choose their favorite music band and leave it public. We could also see that 30% of users who have some data posted on Facebook, have at least one company with which they have been connected.

Also another interesting fact is the difference in the number of men and women, while the number of men is 12 percentage points higher than the number of women. This means that from all mails on the Internet, there are more men that are associated with their email to Facebook. The most probable reason is that men are less conscious of leaving their mail addresses public on Internet.

But the main conclusion to be drawn has been that it has achieved the main objective of the project: to show that we can develop advanced techniques for generating personalized mail that circumvent current spam detection systems.

Clear examples of this are the results shown in the results section. In the first experiment, we can see that only the 0.41% of users have bitten the bait. Whereas in the second 7.62% of the users have entered to the project website. The second result rate is more than 18 times higher than the first one.

We can see that it is not a large number of people, but as a steady stream of visitors, which means that personalized emails reach their destination. Then, once the message is on the user's email inbox, it depends on each person's behavior to click on the link that is sent in the mail. This shows that spam is not blocked as it's customization have not been detected.

Acknowledgments. This work has been partially funded by the Basque Department of Education, Language policy and Culture under the project SocialSPAM (PI_2014_1_102).

References

1. Kanich, C., Kreibich, C., Levchenko, K., Enright, B., Voelker, G.M., Paxson, V., Savage, S.: Spamalytics: an empirical analysis of spam marketing conversion. In: Proceedings of the 15th ACM conference on Computer and communications security. CCS '08, New York, NY, USA, ACM (2008) 3–14
2. KasperskyLab: Spam and phishing in q3. `http://www.kaspersky.com/about/news/spam/2014/iPhones-and-Ice-Buckets-Used-to-Promote-Junk-Mailings`
3. Facebook: Facebook: Newsroom. `http://newsroom.fb.com/company-info/`
4. Bonneau, J., Anderson, J., Danezis, G.: Prying data out of a social network. Social Network Analysis and Mining, International Conference on Advances in (2009) 249–254
5. Balduzzi, M., Platzer, C., Holz, T., Kirda, E., Balzarotti, D., Kruegel, C.: Abusing social networks for automated user profiling. In: Proceedings of the 13th international conference on Recent advances in intrusion detection. RAID'10, Berlin, Heidelberg, Springer-Verlag (2010) 422–441
6. Polakis, I., Kontaxis, G., Antonatos, S., Gessiou, E., Petsas, T., Markatos, E.P.: Using social networks to harvest email addresses. In: Proceedings of the 9th annual ACM workshop on Privacy in the electronic society. WPES '10, New York, NY, USA, ACM (2010) 11–20
7. Alexa Internet, I.: Alexa top 500 global sites. `http://www.alexa.com/topsites`
8. Stringhini, G., Kruegel, C., Vigna, G.: Detecting spammers on social networks. In: Proceedings of the 26th Annual Computer Security Applications Conference. ACSAC '10, New York, NY, USA, ACM (2010) 1–9
9. Gao, H., Hu, J., Wilson, C., Li, Z., Chen, Y., Zhao, B.Y.: Detecting and characterizing social spam campaigns. In: Proceedings of the 17th ACM conference on Computer and communications security. CCS '10, New York, NY, USA, ACM (2010) 681–683
10. Jakobsson, M., Johnson, N., Finn, P.: Why and how to perform fraud experiments. IEEE Security and Privacy **6**(2) (2008) 66–68
11. Jakobsson, M., Ratkiewicz, J.: Designing ethical phishing experiments: a study of (ROT13) rOnl query features. In: WWW '06: Proceedings of the 15th international conference on World Wide Web, New York, NY, USA, ACM (2006) 513–522