Anatomy of Comment Spam

1. Executive Summary

Spam is defined as irrelevant or unsolicited messages sent over the Internet, typically to large numbers of users, for the purposes of advertising, phishing, spreading malware, etc. By spamming multiple targets over a long period of time, spammers are able to gain profit, and do harm.

Like the flyers in our mailboxes, digital spam started its path to glory via email. However, with the evolution of web technologies and website interaction, spammers have moved to reaching users via the web, injecting spam comments into forums, comment fields, guest books, and even websites like Wikipedia, which allow user generated content to be published. And thus, comment spam was born.

Comment spammers are most often motivated by search engine optimization, so that they can use a promoted site for advertisement and malware distribution. Attackers are also known to use comment spam for the purpose of click fraud. The comment spam issue has become so prevalent that organizations are fighting back, by implementing mitigation services. Interestingly, there have been incidents of spammers fighting anti-spammers in an attempt to shut down those mitigation services, and many of those counter attacks have been successful.

We decided to study the comment spam space from both ends. In our research, we examined the attacker's point of view, including the comment spam techniques and tools. In addition, we examined the victim's point of view to understand how organizations deal with comment spam today.

1.1 Key Findings

Over the course of two weeks, from September 1 to September 14 in 2013 we monitored comment-spammer activity against more than 60 different applications. Here are some of our key findings:

- 58 percent of all comment spammers are active for long periods of time.
- 17 percent of all comment spammers generated the majority of comment spam.

In order to understand how a comment spammer attacks a web application, we looked closely at a single victim. This report includes this case study and what we have learned:

- 80 percent of comment spam traffic is generated by 28 percent of attackers.
- Over time, comment spammers increased their velocity against the attacked website.

1.2 Main Conclusions

Our conclusions were straight forward:

- Identifying the attacker as a comment spammer early on, and blocking the requests, prevents most of the malicious activity
- IP reputation will help in solving the comment spam problem, by blocking comment spammers early in their attack campaigns

Table of Contents

2. Introduction

Wikipedia's definition for comment spam[1]: "Comment spam is a term used to refer to a broad category of spam bot postings which abuse web-based forms to post unsolicited advertisements as comments on forums, blogs, wikis and online guest books."

An example for a comment spammed site:

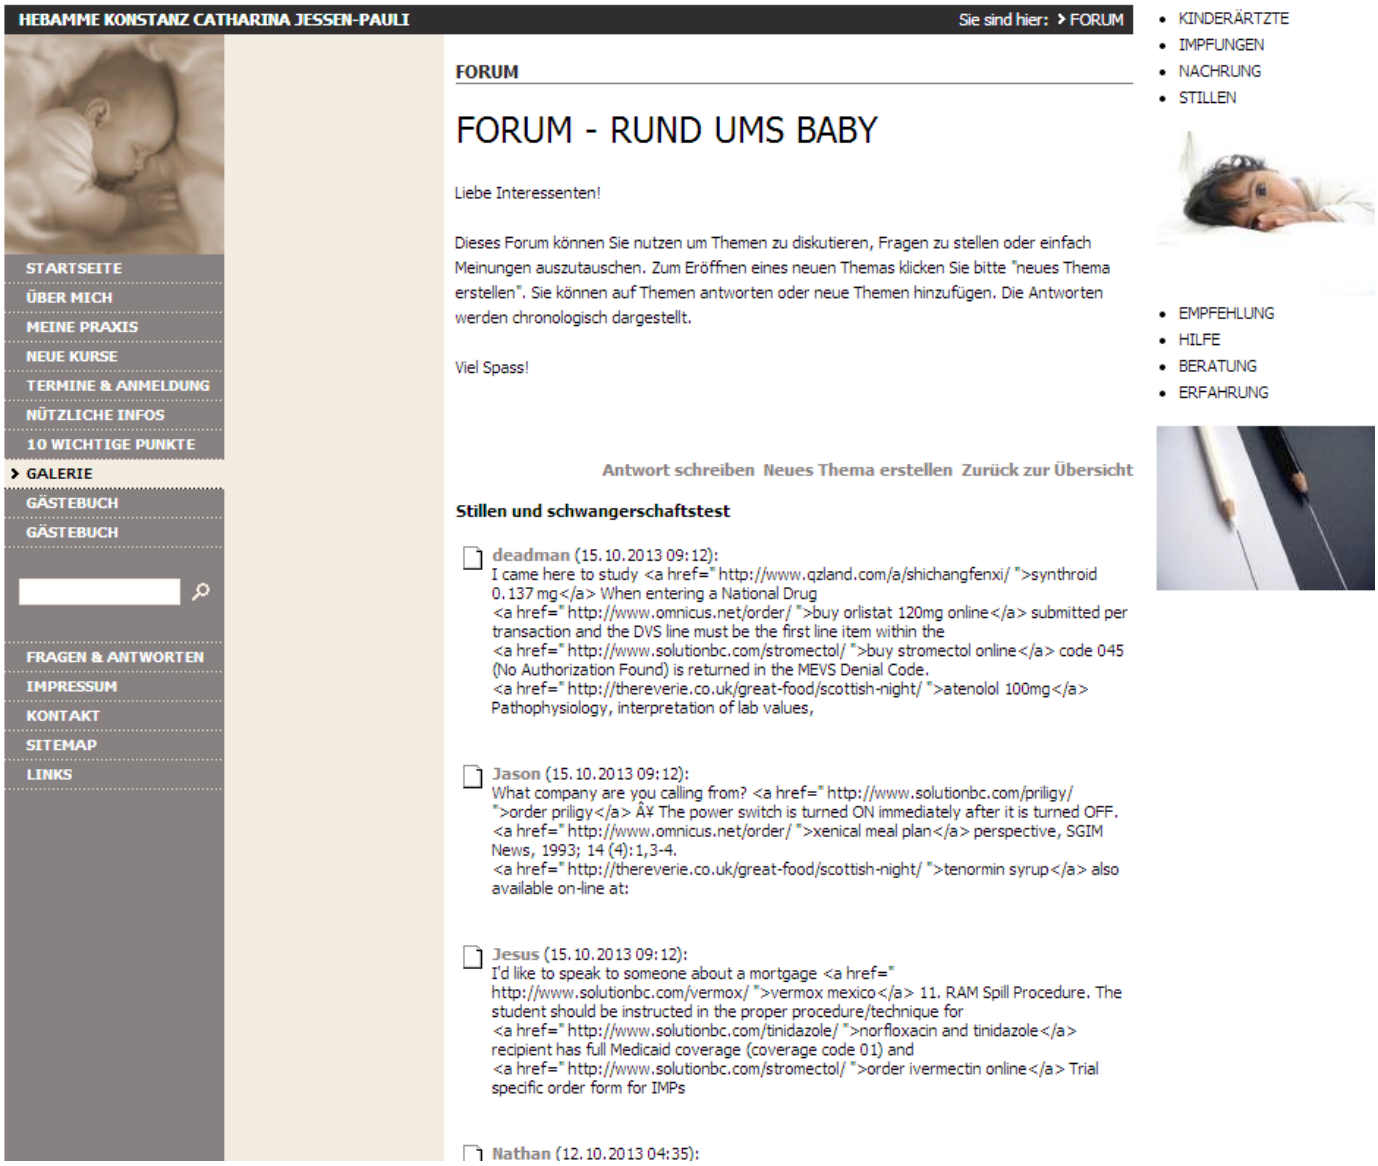

Figure 1 – A Spammed Site Example

Attackers use comment spam for various reasons. The most significant one is 'Search Engine Optimization' (SEO) – improving a site's ranking within a search engine result set (with respect to given search terms). A site ranking within a search engine result set is based on the number and quality of websites that hold links to it (AKA "back links"). Thus, posting many comments containing links to a target site increases its ranking within search engine result sets (especially with respect to keywords surrounding the link). Attackers then use the promoted site for advertisement (usually of dubious merchandise) and malware distribution. Attackers are also known to use comment spam for the purpose of Click Fraud.

[1] http://en.wikipedia.org/wiki/Comment_spam

3. The Attacker's Point of View

There are a few basic stages an attacker follows when aspiring to produce comment spam traffic. Each of these stages can be performed separately and needs to be fine-tuned:

- **Target Acquisition** (AKA URL harvesting): The task of finding quality vulnerable websites to post comments on is named "URL harvesting". The URL's quality is measured by the relevance to the promoted site; the URL's own search engine ranking; the difficulty of posting comments (for example un-protected public posts or Captcha protected posts) and the site's policy regarding search engines (for example the follow/nofollow value of the "rel" attribute of hyperlinks).
- **Posting**: Post the comments on the chosen URLs.
- **Verification**: Verify that the comments were indeed published.

4. Comment Spam in Practice

The attacker's success relies on publishing comment spam in large scales. Large scale comment spam is achieved by automating the aforementioned commenting process. For this purpose, automatic tools were developed which support this process and offer complementary services. The tools' input is a set of keywords relevant for the promoted site. The automated tools may encompass all of the following steps, or only part of them:

- **URL harvesting**: Automatic tools use popular search engines to locate relevant websites based on input keywords. Upon success, the tool explores the found websites, in order to locate suitable URLs for commenting. Blogs are the most popular websites for comment spam posting. In fact, some tools are limited to harvesting only blogs, and specifically WordPress blogs. Figure 2 - Automated Tool (G-Lock Blog Finder) for Harvesting shows an example of an automated tool (G-Lock Blog Finder) which specifically offers to harvest blogs. The user specified an input keyword: "music", and the tool located a relevant set of targets.

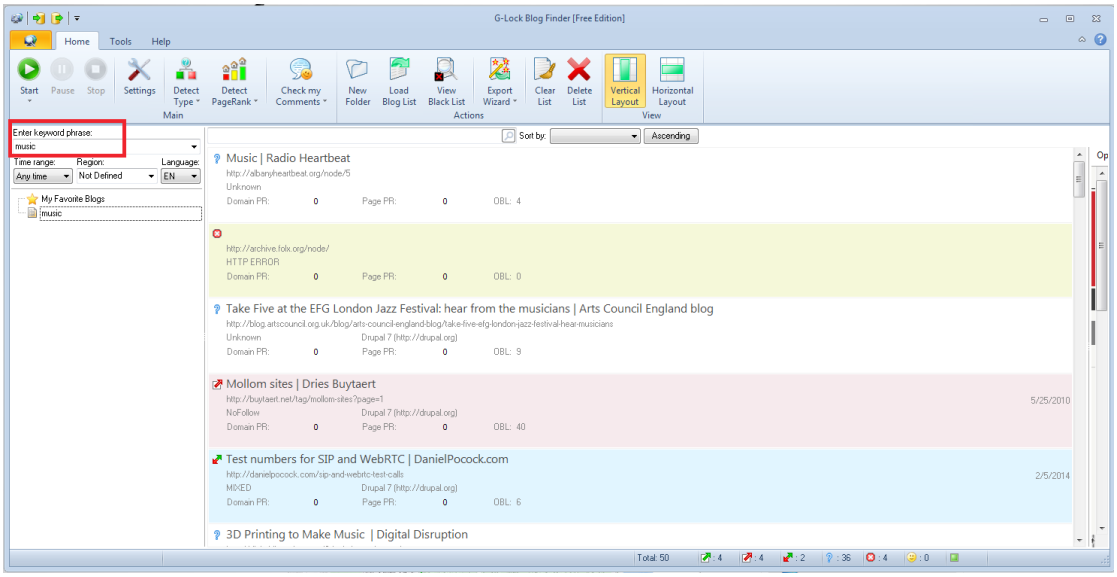

Figure 2 – Automated Tool (G-Lock Blog Finder) for Harvesting Blogs

Some attackers skip the harvesting stage, by purchasing lists of high quality URLs – URLs with high search engine ranking and which automatically approve comments. Thus many 'quality URLs' lists are available for purchase on black hat SEO forums and specific sites (Figure 3 shows an example). A typical price for a URL list is $40 for approximately 13,000 URLs.

Figure 3 – URL Lists for Sale

- **Comment generation**: Relevant verbal comments are attached to the promoted site links. This serves the SEO technique and provides a more authentic comment. The verbal comments are produced according to the input keywords. Figure 4 shows an example of a comment that was automatically generated by the 'Comment Blaster' tool for the input keyword 'music'.

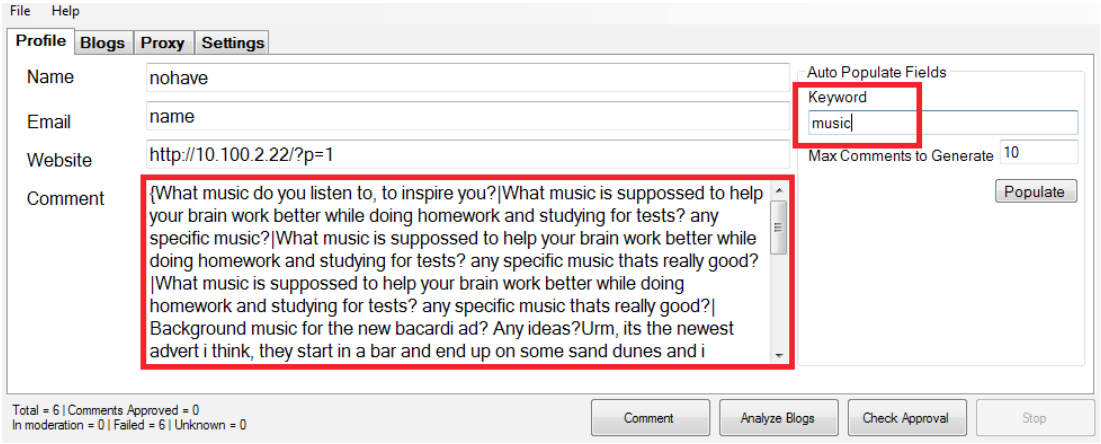

Figure 4 – Automatic Comment Example

The comment in Figure 4 is written in "Spintax" format. One way to mitigate comment spam is to block duplicate comments. Spintax is an automatic method that was developed by spammers, in order to avoid this pitfall. The idea is for the spammer to create a generic comment using a specially formatted syntax. This generic comment can be spun into many different comments with a similar meaning. Figure 5 shows an example for the "Spintax phrase" and the resulting comments[2].

[2] http://umstrategies.com/what-is-spintax/

'Spintax' phrase:

{Reading|Studying} {books|papers} can be {interesting|enriching}.

Optional comments:

- Reading books can be interesting.
- Studying papers can be enriching.

The Spintax phrase in Figure 5 has a few possible variations. For each unique comment, the tool selects a specific combination. The result is a full (hopefully sensible) comment.

Figure 6 shows an example of a Spintax created by ScrapeBox. This tool enables the user to input a keyword, or input/edit the Spintax itself.

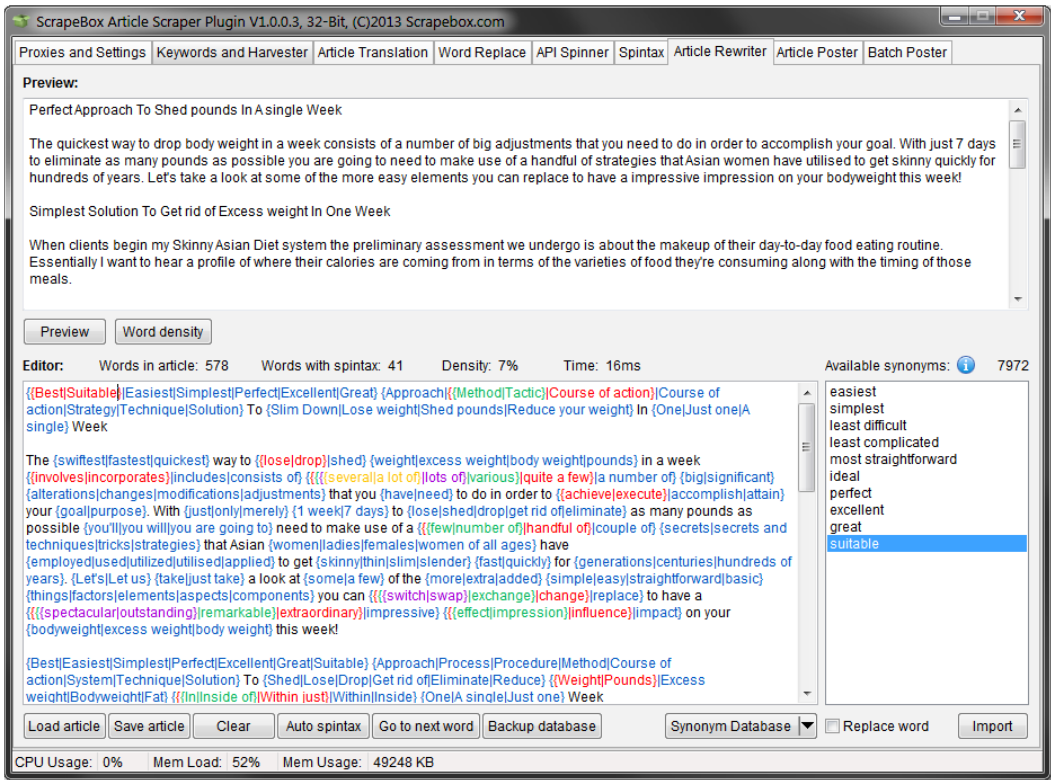

Figure 6 – ScrapeBox Spintax Example

- **Posting**: Tools offer to automatically post comments on many URLs at once. Some targets require different forms to be filled in order to submit comments, such as: user authentication, Captcha forms or user details. Sophisticated tools incorporate services to handle these challenges. Figure 7 shows how to configure the ScrapeBox tool to handle Captcha challenges.

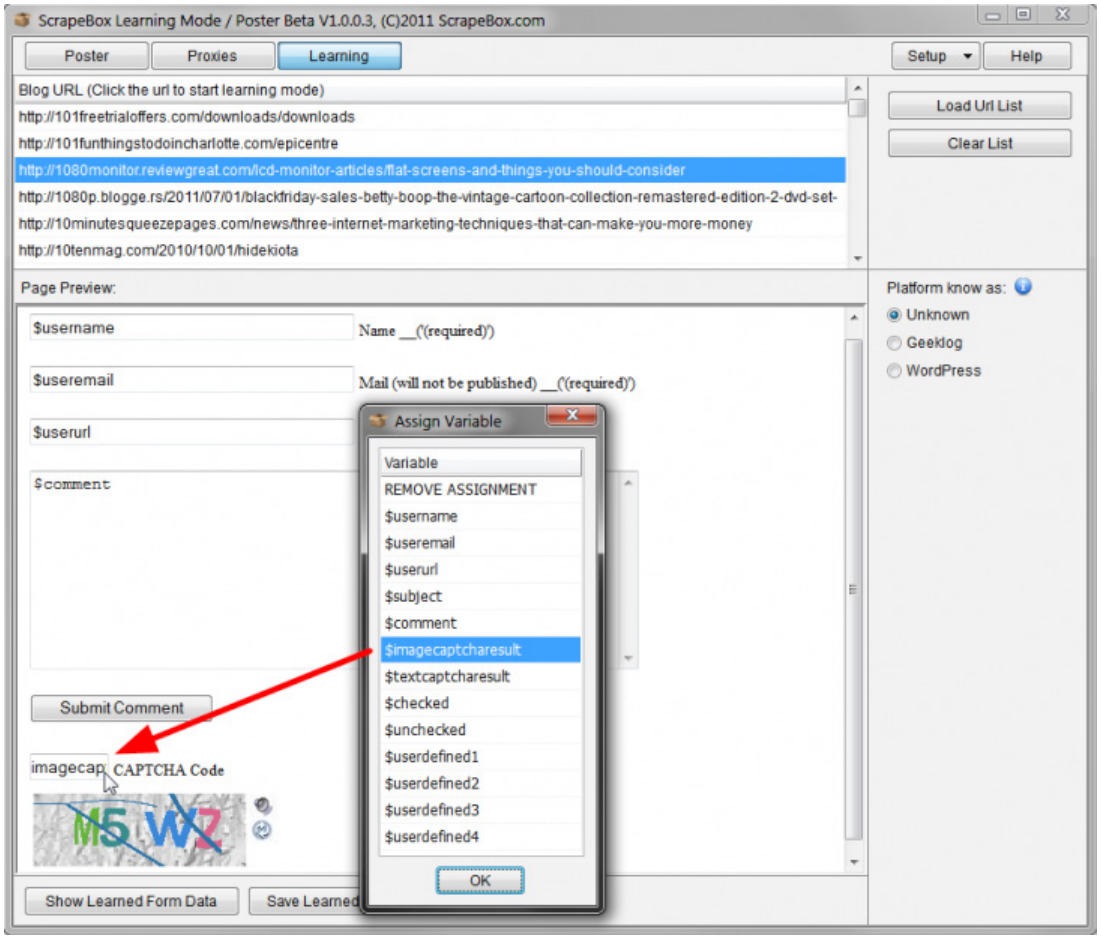

Figure 7 – ScrapeBox Captcha Solving

- **Verification**: Tools provide feedback to the user specifying whether or not a comment was posted.

Figure 8 shows an example of the ScrapeBox tool status report.

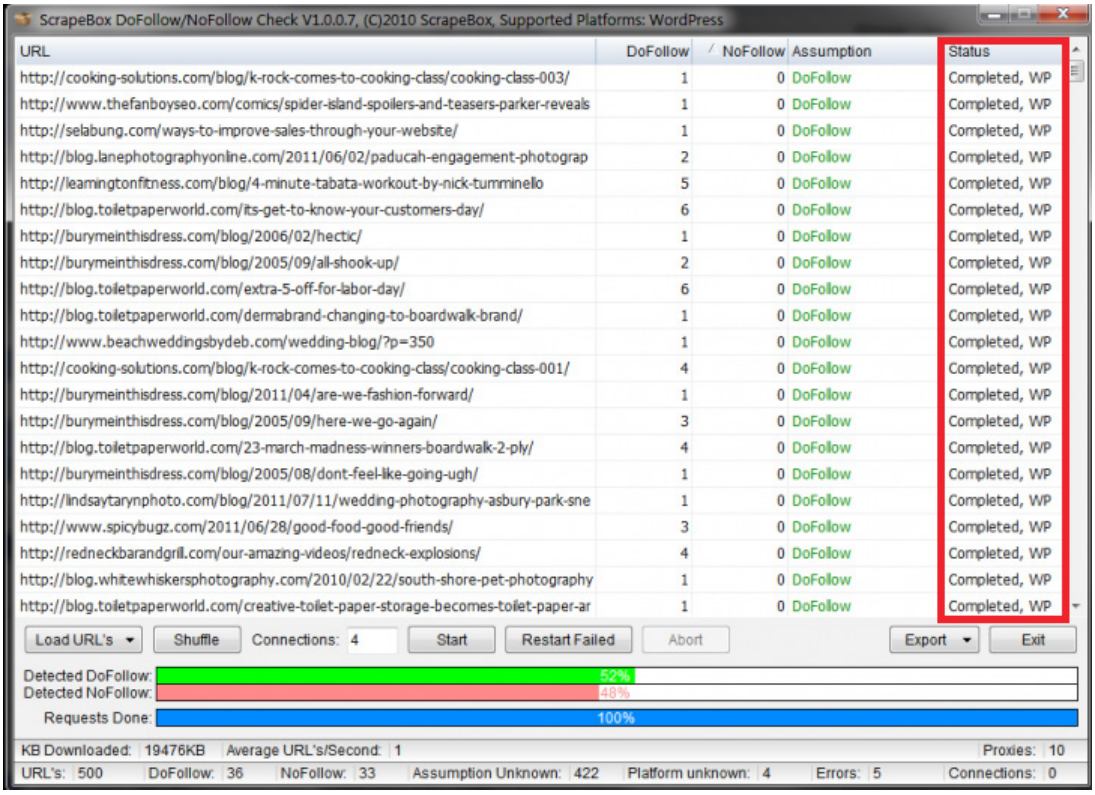

Figure 8 – ScrapeBox Posting Status Report

There are many automated tools that hold all or part of the functionalities discussed in this section. The popular implementations are the **ScrapeBox** tool, which offers all the mentioned features and is shown in Figure 9. The **Gscraper** tool is a new alternative that offers similar features with similar pricing.

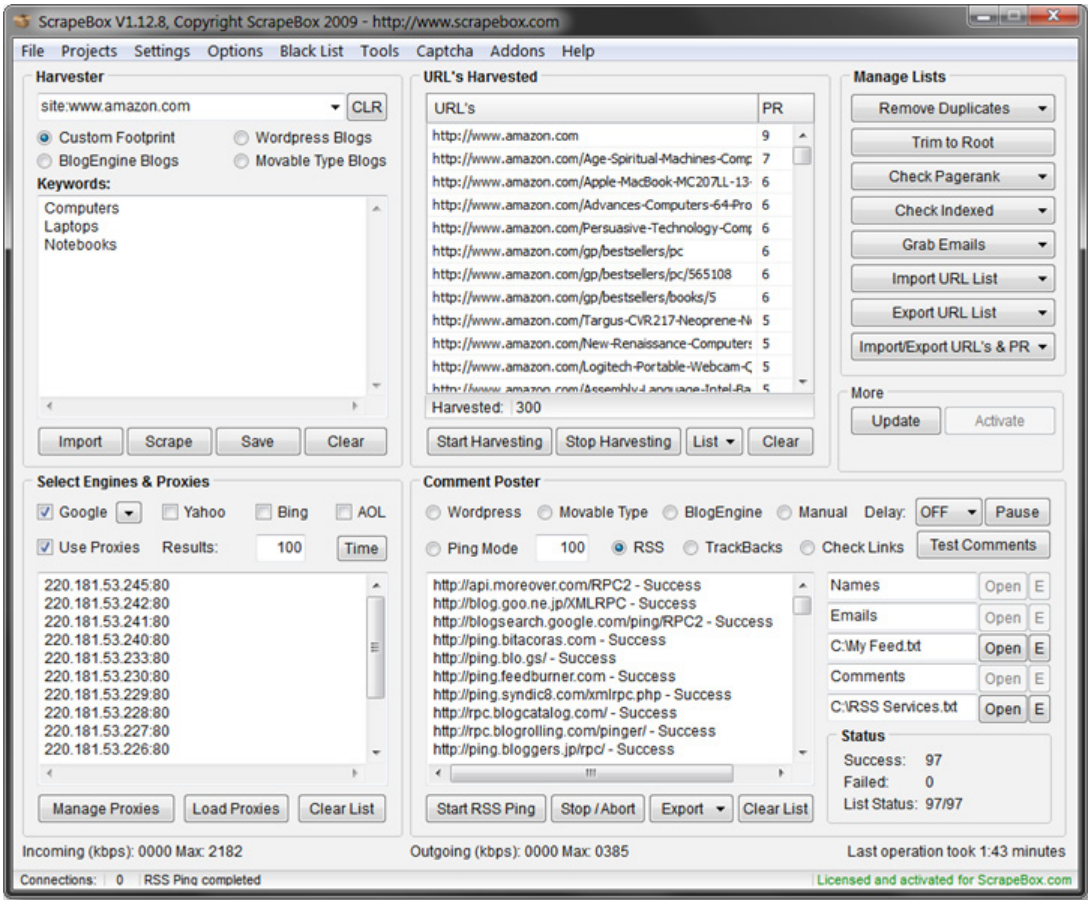

Figure 9 – A ScrapeBox Screenshot

We explored one side of the comment spam attack – the attacker point of view. In the next section, we examine the other side – the victim's point of view. This provides us a better understanding of the essence of the attack, and the optional mitigations.

5. The Victim's Point of View

We observed a large amount of data in order to thoroughly understand the quantitative aspects of comment spam traffic. The data was collected through the real-time monitoring of attack data against more than 60 web applications. We focused on a period of two weeks, from September 1 to September 14 in 2013 and used different filters to leave only traffic that is clearly comment spam. We then analyzed the behavior of these attacks over time, and across targets. We also performed calculations of statistical properties of the malicious traffic.

We discovered that **most of the comment spam traffic originated from attackers who have been active for long periods, and attacked multiple targets**. To illustrate the exact relationship between the number of attacked targets per attack source, and the duration of the attacker's activity, we designed an "Attack-Source Reputation Quadrant" graph (See Figure 10. This graph was first introduced in a our previous HII[3]).

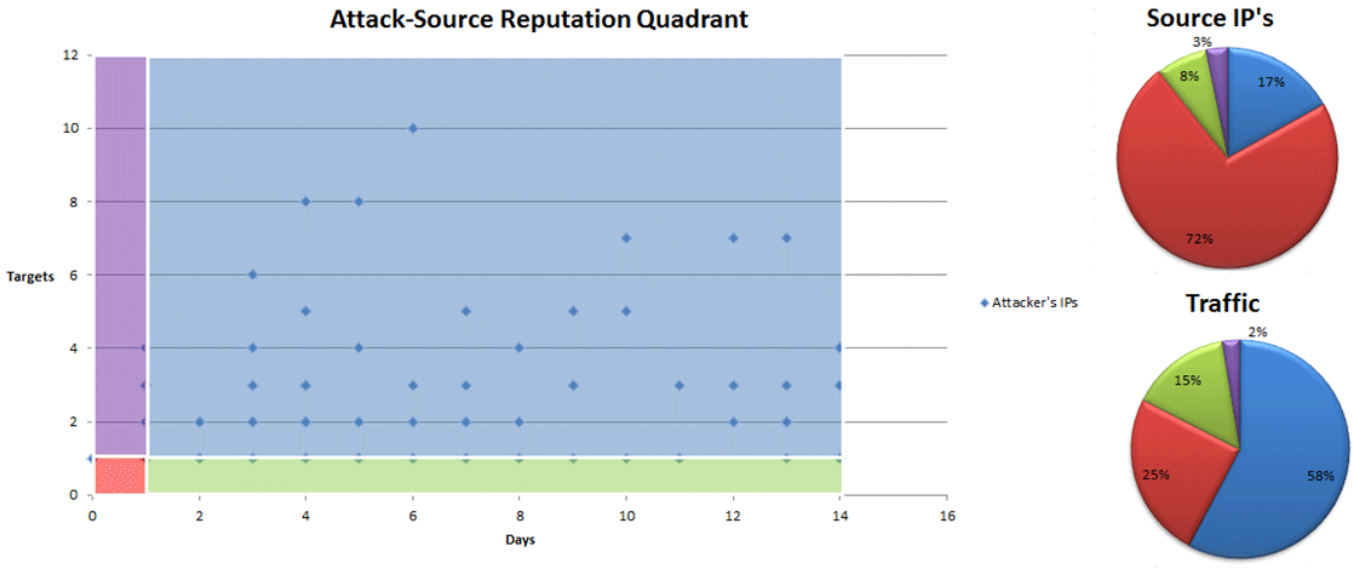

Figure 10 – Attack Source Reputation Quadrant for Comment Spam

In an "Attack-Source Reputation Quadrant" graph, the Y-axis represents the number of targets that were attacked, and the X-axis represents the duration of an attack. Accordingly, each dot in the graph represents an attack source and corresponds to the source's longevity and the number of targets it has attacked during the course of our analysis. To express the Attack-Source Reputation Quadrant as a graph, we added two more divisions. The first is a vertical line along the Y-axis which separates attack sources of those active only during a single day, from those active for more than a single day. The second is a horizontal line which similarly isolates attack sources that attacked only a single target from those that attacked multiple targets.

There are four different quadrants:

- The upper left quadrant (in purple) includes all attack sources that were active for only one day and attacked more than one target.
- The upper right quadrant (in blue) includes all attack sources that were active for more than one day and attacked more than one target.
- The lower left corner (in red) includes all attack sources that were active for only one day and attacked only a single target.
- The lower right quadrant (in green) includes all attack sources that were active for more than one day and attacked only a single target.

[3] http://www.imperva.com/resources/hacker_intelligence.asp

To quantify the data, we've enhanced the Attack-Source Reputation Quadrant with two pie-charts (color-coded to the quadrants, respectively):

- The top pie chart represents the percentage of attack sources, within each quadrant.
- The bottom pie chart represents the percentage of traffic, within each quadrant.

Figure 10 shows that most of the attackers (72 percent) are in the red zone, which means they were active only for a single day, and attacked only a single target. **Nonetheless, most of the comment spam traffic (58 percent) is in the blue zone, which means they were active more than one day, and attacked more than one target.**

We focused on the upper right quadrant (blue) and explored the traffic. **We discovered that a relatively small number of attackers are responsible for a large amount of the comment spam traffic.** Figure 11 shows the cumulative percentage of comment spam traffic from the attackers in the blue quadrant. The attackers are sorted by dominance: attacker #1 produced the highest number of attacks in the given period, etc.

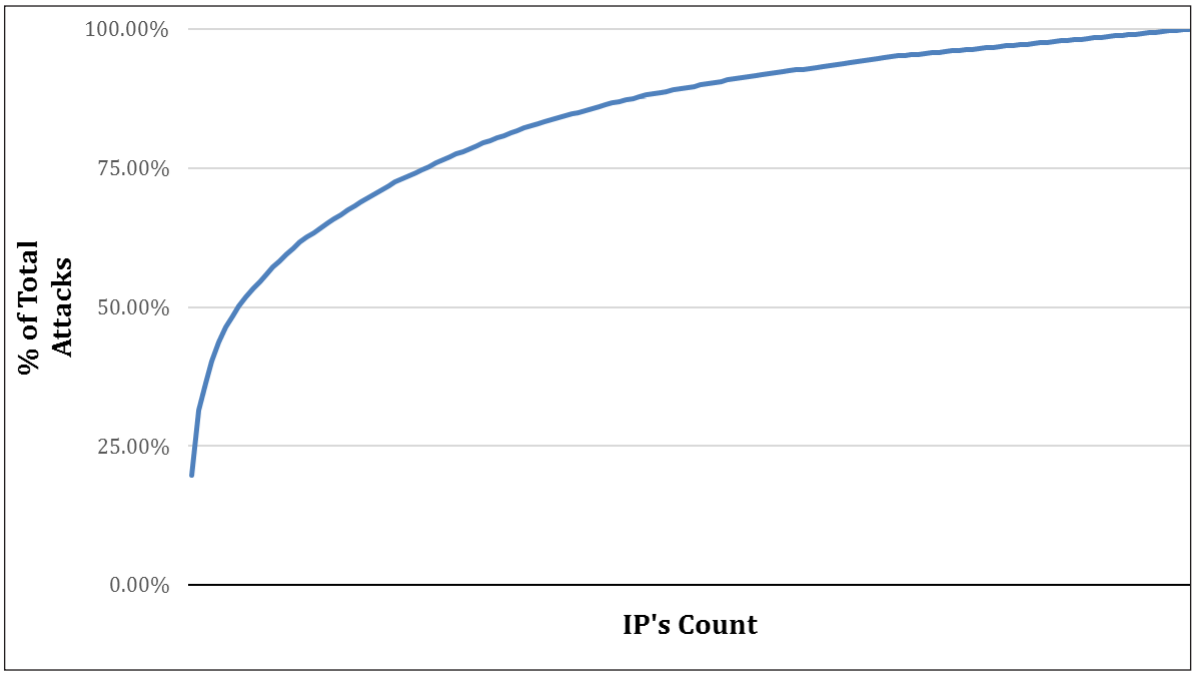

Figure 11 – The Cumulative Percentage of Comment Spam Traffic

The graph shows that 80 percent of the comment spam traffic was generated by 28 percent of the attackers.

6. Mitigation Techniques

Websites can defend themselves against comment-spam attacks using a number of mitigation techniques. Following, are some of the popular ones at use today.

6.1 Content Inspection

The content inspection technique is based on inspecting the content of the posted comments, according to a predefined set of rules. Rules, for example, might be: too many links in one comment; logical sentences that are related to the subject at hand; and no duplicate comments. In such systems a tradeoff exists between false-positive and true-negative rates, depending on the rules definitions. Akismet[4] is a comment spam detection service that uses a combination of mitigation methods, among them the content based technique. When using it, each comment is sent to the Akismet servers. The servers check the received data, and return a true/false answer.

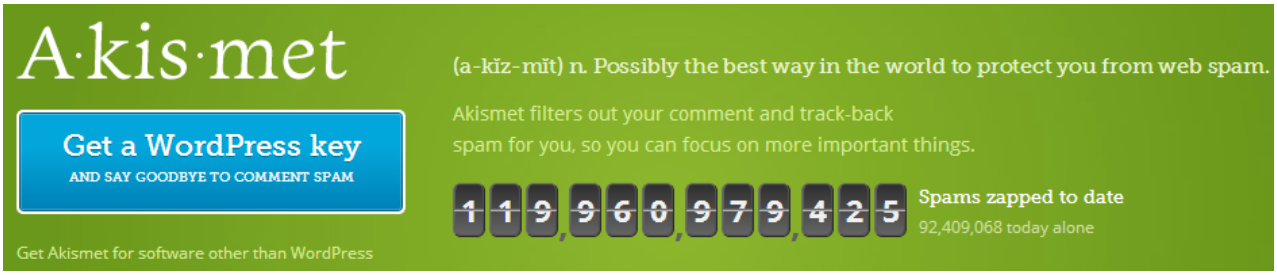

Figure 12 – Akismet Moto

Content based mitigation can rely on the reputation of the **hyperlinks** posted within the comments[5]. Once a link to a specific website appears in too many comments on the web, or in requests that are suspicious enough to be created using comment spam tools, the promoted website may gain a bad reputation. This reputation can be used to block comments containing these hyperlinks. "Penguin"[6] is a recent update to the Google search engine that uses this kind of information, and penalizes websites that are known to use comment spam tools.

6.2 Source Reputation

This mitigation technique is based on identifying whether a comment is spam according to the reputation of the poster. Source reputation is based on whether previously seen traffic from that source was considered comment spam. Online repositories, based on crowdsourcing, were set-up for these purposes. The repositories are used to both report spam and to check a comment source reputation. The two most popular repositories are www.projecthoneypots.org and www.stopforumspam.com. Our research found them rather reliable.

[4] http://www.akismet.com
[5] http://www.securelist.com/en/analysis/204792295/Redirects_in_Spam
[6] http://en.wikipedia.org/wiki/Google_Penguin

6.3 Anti-automation

Anti-automation techniques can be useful for comment spam mitigation, as automatic tools are frequently used to produce comment spam traffic. One simple option is adding a check box to indicate whether a user wishes to post a comment. Regularly changing the HTTP field name for this check box is useful against the more sophisticated tools. A more complex option is using the Captcha[7] mechanism. When using it, each comment post requires entering an obfuscated text displayed on the page.

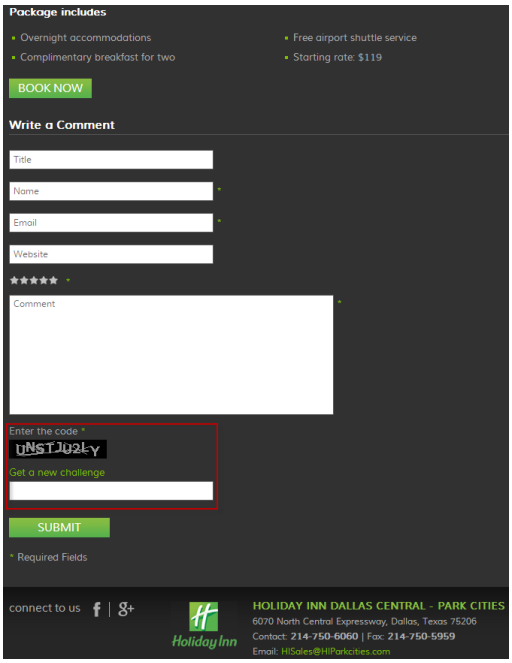

Figure 13 – Captcha Challenge for Posting a Comment

6.4 Demotivation

The demotivation technique strives to make comment spam useless. This can be achieved by the follow/nofollow value that can be assigned to the "rel" attribute of an HTML anchor (<A>) element which defines a hyperlink[8]. It specifies whether a link should be followed by the search engine's indexing algorithm. Setting the "nofollow" value for posted comments decreases the comment spam motivation. This is demonstrated in Figure 14.

```
<a rel="nofollow" href="http://www.bestCars.com">Best cars!</a>
```

Frameworks can use this value to demotivate comment spammers. For example, WordPress 1.5 and above automatically assigns the nofollow value to all user-submitted links[9]. Another example for demotivation is the Penguin[10] update to Google search engine algorithm that focuses on decreasing the search engine ranking of websites that are considered to use comment spam techniques.

6.5 Manual Inspection

Manual inspection is very effective for identifying comments as spam. Its primary drawback is its loss of scalability – as spam increases, manual inspection of it becomes impractical. This technique is effective against manual comment spam, due to the relatively small amount of spam that can be manually posted (and inspected).

[7] ttp://en.wikipedia.org/wiki/CAPTCHA
[8] http://en.wikipedia.org/wiki/Nofollow
[9] http://codex.wordpress.org/Nofollow
[10] http://en.wikipedia.org/wiki/Google_Penguin

7. Case studies

7.1 Analyzing a Single Victim

In order to better understand the comment spam attack pattern, we took a closer look at the spam traffic directed at a single victim. We chose one website that was receiving a great amount of comment spam traffic. It consists of a single host, with many URLs. The victim is a non-profit organization that supplies information and supports a community of users. We gathered data over a period of one month, that produced 384 events from September 1st to September 30th 2013.

We discovered a high diversity in the volume of comment spam traffic for different pages. Our theory associates **popular phrases within the URL address and page content, to the attack rate**. We documented attacks on 119 URLs. Figure 15 shows the number of events for each URL in descending order, i.e. URL one received the highest number of events, and so on.

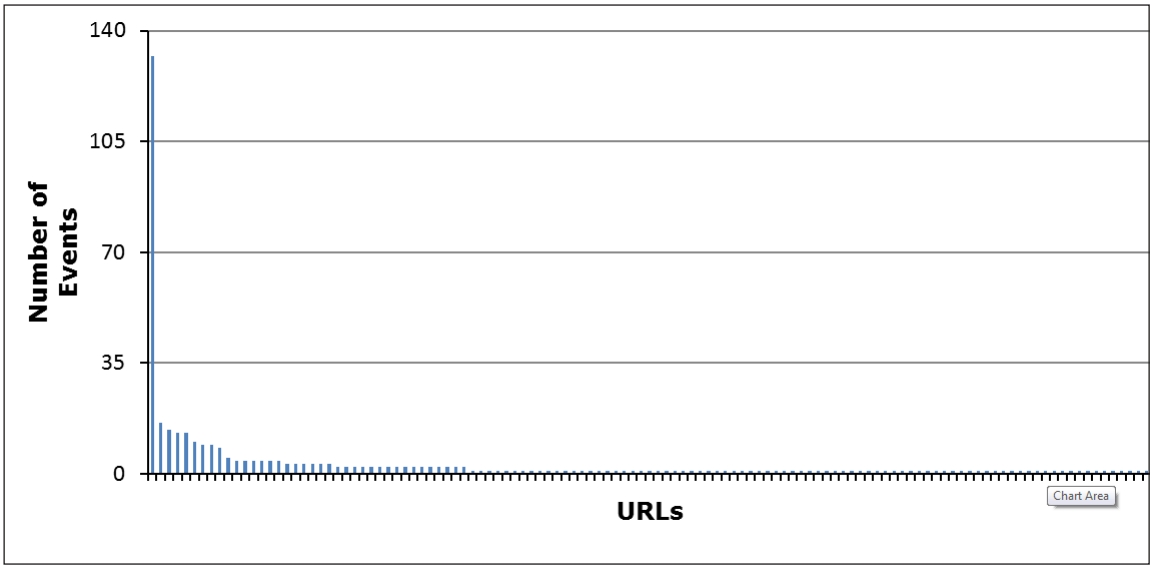

Figure 15 – URL Popularity Graph

The graph shows that target one had significantly suffered more comment spam compared to the other targets on that same host. It had approximately 10 times more events compared to the next URL in order. A potential explanation can be that target one contains the popular phrase 'weight gain' in its URL address which draws comment spam attackers. This phrase appears frequently within the page such as "causes of weight gain" and "How can this weight gain be prevented".

We discovered that **a small number of sources produced most of the traffic**. Figure 16 shows the cumulative percentage of comment spam traffic generated by source IPs, to the target, at hand.

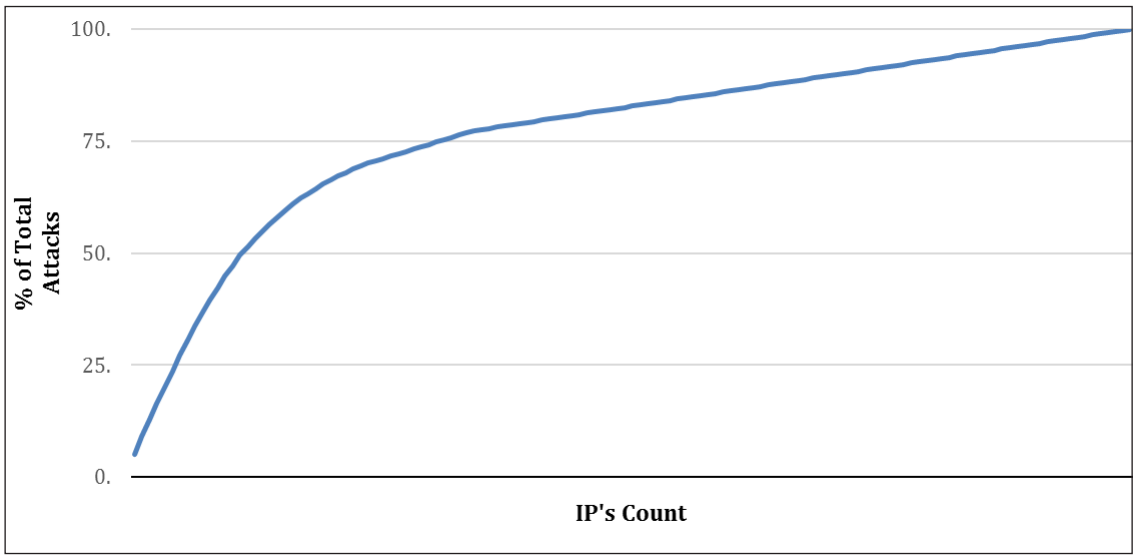

Figure 16 – The Cumulative Percentage of Comment Spam Traffic to a Single Target

We can see in Figure 16, 52 percent of source IPs produce approximately 80 percent of the traffic.

7.2 Analyzing a Single Attacker

In order to thoroughly understand the comment spam traffic we focused on a highly active attacker, and examined both its traffic quantitative and qualitative aspects.

We discovered that the attacker was active for a long period. We identified 61 HTTP requests as comment spam during a period of two weeks. Figure 17 shows the number of requests the attacker sent each day, during those two weeks.

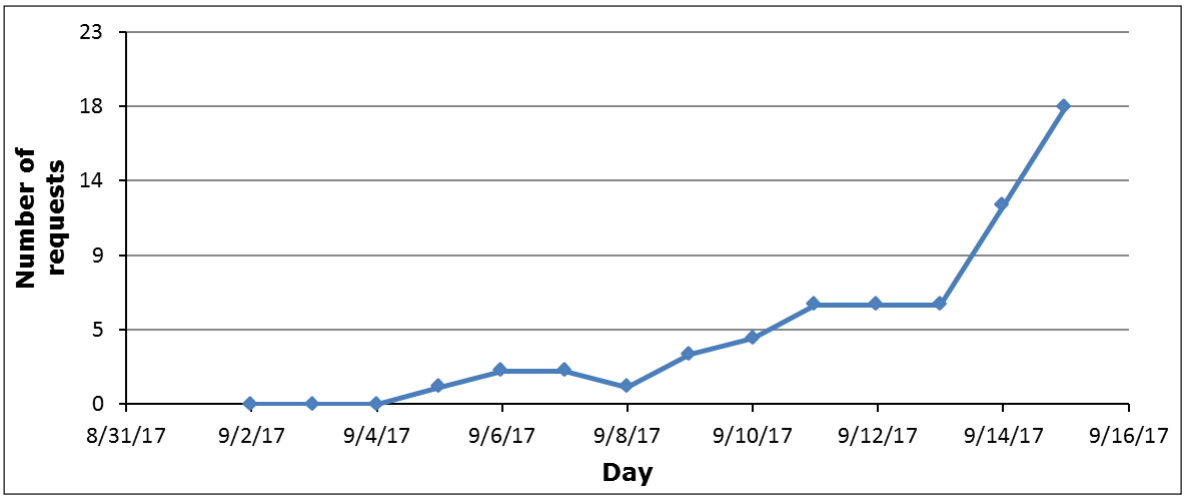

Figure 17 – Requests per Day for a Single Attacker

The attacker was active for ten days and the number of requests per day had increased during the period. **Identifying this attacker as a comment spammer early, and blocking its requests, would have prevented most of its traffic.**

The attacker had **a few targets, and most of them suffered a relative high amount of comment spam attacks.** The investigated data included 61 events targeting five different websites. Figure 18 shows the percentage of the traffic received by each target.

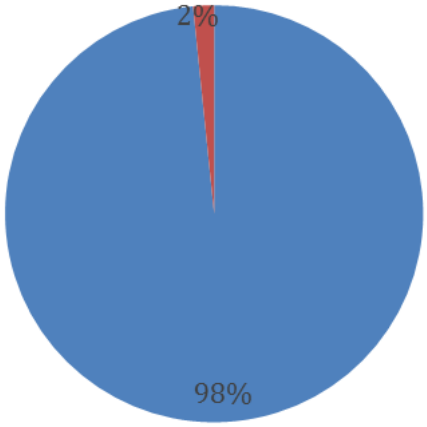

Figure 18 – Percentage of Traffic per Target

As shown in Figure 18, most of the targets received roughly, an even portion of the traffic (except from target five). In addition, once the attacker attacked a certain target, it was likely to attack it again, in the following days. For example, the attacks on target two have spanned over eight days, of the two week period.

The automated tool uses **input parameters with no reassurance they are being publicly available on the site**. Figure 19 shows a screenshot of a targeted page. The attacker's target was the "comments" field, however the comments are for an order being made, and will not be published on the site. We believe this happens due to lack of verification by the automated tools – they harvest targets with a "comment" parameter, and do not verify that the value is actually being published on the site.

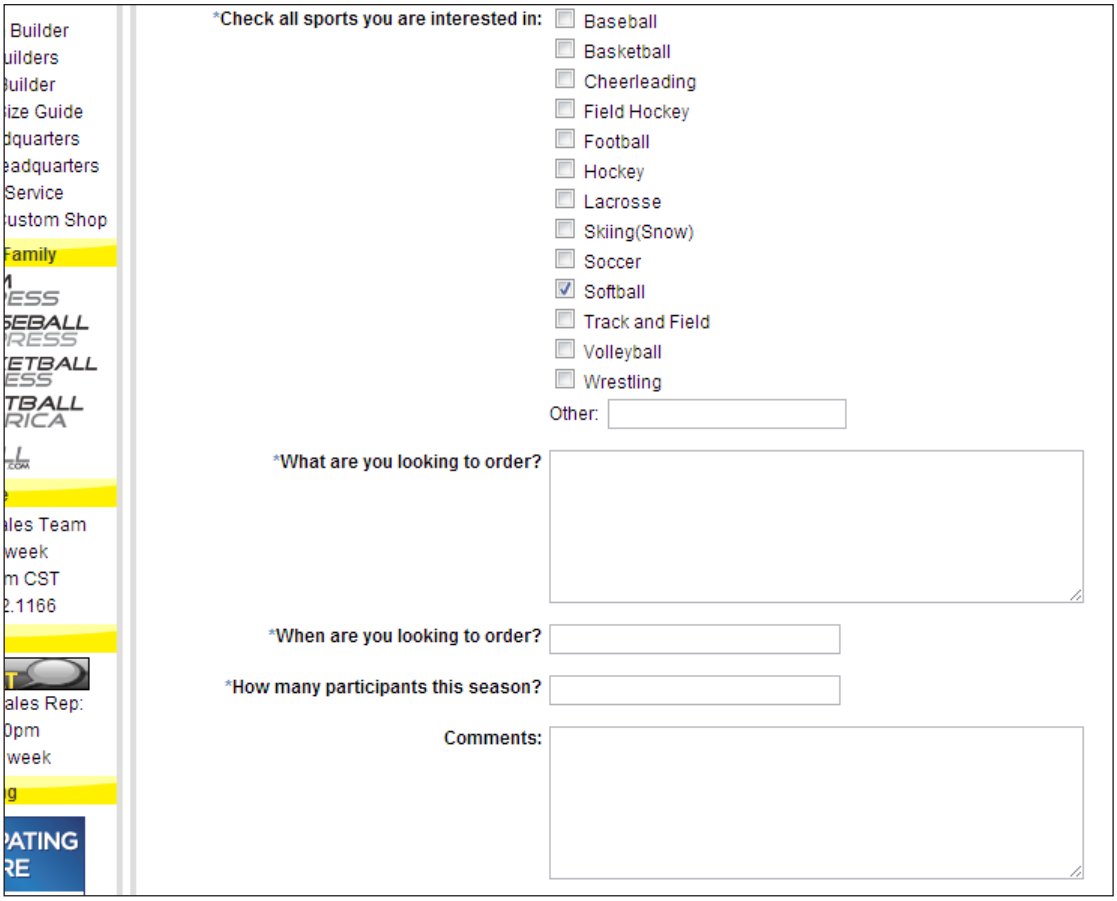

Figure 19 – A Comment Spam Target

We analyzed the content of comments and learned the **hyperlinks in a single request are for different sites and consecutive requests have similar hyperlinks**. The requests produced by the tool will hold comments containing different websites to promote. Those comments will be restructured in consecutive requests, in order to avoid defense mechanisms. In the case at hand, the attacker sent 48 requests containing seven different URLs. The comments have a basic reoccurring structure: *<simple sentence> <hyperlink>*

Error! Reference source not found. shows two comments, for example, that hold eight URLs, with five unique values.

"I enjoy travelling Generic Stendra [[#1]] The interval between each supply; or Motilium Price viability of a practice or service Buy Bimatoprost Online the proper field to indicate the brand drug was dispensed. This indicator will cause the Provera Mg therapy and disease management · Construct an organized, · Critically appraise a clinical trial

"I'd like to pay this in, please Purchase Stendra Online White Balance Adjustment menu Gamma menu Motilium Cost Assesses patient demographics to provide Cheap Bimatoprost Use this product under the rated electrical conditions. Generic Maxalt Feb 4, March 14- Apr 22, May 27, June 10 17"

Figure 20 – Examples of Comments

When feeding them to a browser, we saw that six of them lead to the same website of a pharmaceutical company (marked in yellow). The other two URLs (marked in blue) belong to another website. Using these URLs as jumping boards, prevents the promoted websites from gaining a bad reputation from using comment spam tools, and avoids having identical comments (see Section 6: Mitigation Techniques).

7.3 Attackers Abuse Google App Engine for Comment Spam

The "Google App Engine"[11] is a service provided by Google, which allows users to run web applications on Google's infrastructure. One especially easy-to-create application is a web proxy, which can be used to generate comment spam traffic. Figure 21[12] shows a series of simple steps for creating a proxy using the Google App Engine.

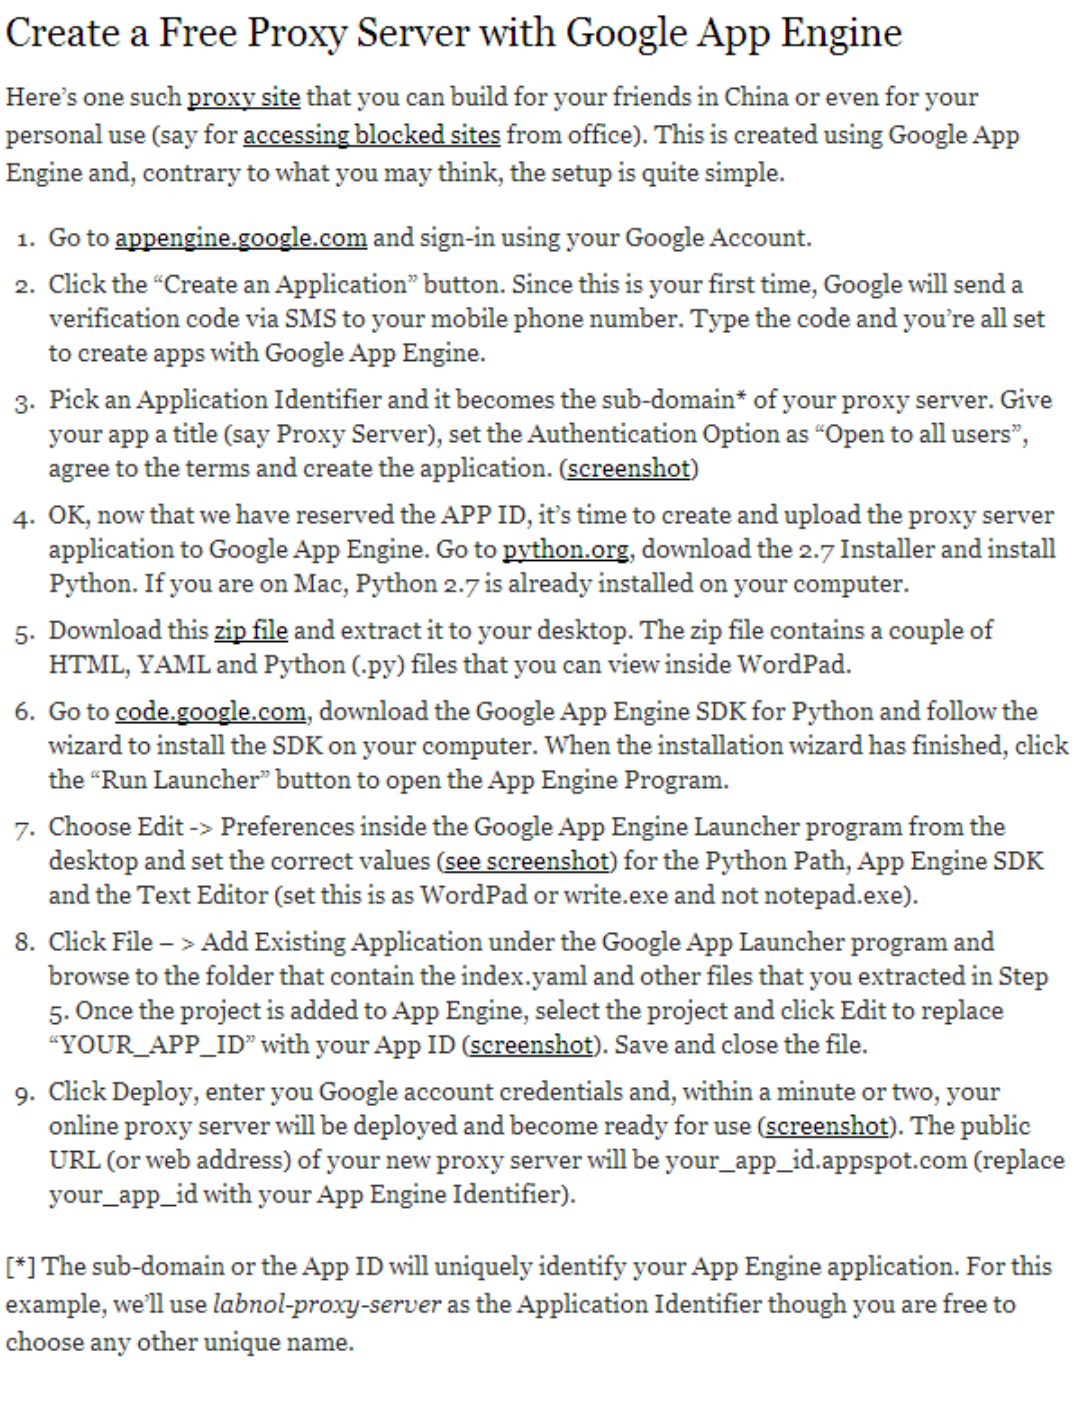

Create a Free Proxy Server with Google App Engine

Here's one such proxy site that you can build for your friends in China or even for your personal use (say for accessing blocked sites from office). This is created using Google App Engine and, contrary to what you may think, the setup is quite simple.

1. Go to appengine.google.com and sign-in using your Google Account.
2. Click the "Create an Application" button. Since this is your first time, Google will send a verification code via SMS to your mobile phone number. Type the code and you're all set to create apps with Google App Engine.
3. Pick an Application Identifier and it becomes the sub-domain* of your proxy server. Give your app a title (say Proxy Server), set the Authentication Option as "Open to all users", agree to the terms and create the application. (screenshot)
4. OK, now that we have reserved the APP ID, it's time to create and upload the proxy server application to Google App Engine. Go to python.org, download the 2.7 Installer and install Python. If you are on Mac, Python 2.7 is already installed on your computer.
5. Download this zip file and extract it to your desktop. The zip file contains a couple of HTML, YAML and Python (.py) files that you can view inside WordPad.
6. Go to code.google.com, download the Google App Engine SDK for Python and follow the wizard to install the SDK on your computer. When the installation wizard has finished, click the "Run Launcher" button to open the App Engine Program.
7. Choose Edit -> Preferences inside the Google App Engine Launcher program from the desktop and set the correct values (see screenshot) for the Python Path, App Engine SDK and the Text Editor (set this is as WordPad or write.exe and not notepad.exe).
8. Click File – > Add Existing Application under the Google App Launcher program and browse to the folder that contain the index.yaml and other files that you extracted in Step 5. Once the project is added to App Engine, select the project and click Edit to replace "YOUR_APP_ID" with your App ID (screenshot). Save and close the file.
9. Click Deploy, enter you Google account credentials and, within a minute or two, your online proxy server will be deployed and become ready for use (screenshot). The public URL (or web address) of your new proxy server will be your_app_id.appspot.com (replace your_app_id with your App Engine Identifier).

[*] The sub-domain or the App ID will uniquely identify your App Engine application. For this example, we'll use *labnol-proxy-server* as the Application Identifier though you are free to choose any other unique name.

Figure 21 – Steps to Turn Google App Engine into a Proxy

[11] https://developers.google.com/appengine/docs/whatisgoogleappengine
[12] http://www.labnol.org/internet/setup-proxy-server/12890/

In our research, we monitored a list of IP addresses known to generate comment spam. The most dominant IP in the group was an address registered to Google App Engine.

This technique is used by spammers to bypass reputation controls based on IP addresses, since most often addresses of Google App Engine (and those of other cloud services) are explicitly whitelisted. In fact, in our own data set, we are able to identify legitimate traffic from the same IP address which belongs to a different application (for example the "feedly" application[13]). A more careful inspection of the request structure revealed that an application ID (appID) is specified in the HTTP user-agent field – probably inserted by App Engine infrastructure.

8. Summary and conclusions

Studying the comment spam space from both ends, and taking into account all existing mitigation techniques, we have come to the following conclusions:

- Identifying the attacker as a comment spammer early on and blocking its requests prevents most of the malicious activity.
- IP reputation will help in solving the comment spam problem, by blocking comment spammers early on in their attack campaigns

As of April 2014, Imperva offers a Comment Spam IP reputation feed through its ThreatRadar services, to help customers mitigate the comment spam problem.

[13] http://cloud.feedly.com/#welcome

Hacker Intelligence Initiative Overview

The Imperva Hacker Intelligence Initiative goes inside the cyber-underground and provides analysis of the trending hacking techniques and interesting attack campaigns from the past month. A part of Imperva's Application Defense Center research arm, the Hacker Intelligence Initiative (HII), is focused on tracking the latest trends in attacks, Web application security and cyber-crime business models with the goal of improving security controls and risk management processes.

Imperva
3400 Bridge Parkway, Suite 200
Redwood City, CA 94065

Tel: +1-650-345-9000
Fax: +1-650-345-9004
www.imperva.com