Analysis on the Effectiveness of Safe Browsing Services

December 7, 2010

Frankie Angai, Calvin Ching, Isaiah Ng, Cameron Smith
frankiea@ieee.org, calvin.k.w.ching@gmail.com, isaiahng@gmail.com, kamrn@interchange.ubc.ca

Department of Electrical and Computer Engineering
University of British Columbia
Vancouver, Canada

***Abstract*—This paper presents an analysis on the effectiveness of Safe Browsing Services (SBS), such as those provided by Norton, McAfee and Google. SBSs help identify and warn end users of potentially malicious websites through popups, splash pages, or other notifications. A website is typically considered to be malicious if the site hosts or links to malware content. Unfortunately, SBSs implement and maintain their own malware detection system and algorithm, resulting in major discrepancies between different services. Furthermore, due to the lack of a universal safe-site policy, each service flags websites as malicious based on different criteria. As a result of such diverse criteria, many legitimate websites have been identified as malicious simply from visitors posting links or images from malicious websites in comments or forum posts. Through an analysis and user study, this paper recommends numerous techniques and methods as countermeasures that SBSs may implement, and also provides a prototype application in the form of a Google Chrome Extension to demonstrate the suggested techniques.**

I. Introduction

OVER the past 50 years, the Internet has provided a platform to deliver a wide range of information and services to over billions of users from around the world. However, it has also provided an outlet for people with malicious intent to steal sensitive information from unknowing users. Phishing and cross-site scripting attacks have become so common and relevant that a tremendous amount of effort has been placed on ensuring a safe browsing experience for regular users [1]. Software security companies like Symantec (Norton) and McAfee, as well as certain search engines, provide Safe Browsing Services (SBS) to protect users from entering sites that may contain malicious content [2]. These services, unfortunately, have various limitations.

Firstly, each of these services compiles and manages their own database of malicious sites. As a result, each service contains a significantly different database of sites than other services, leading to discrepancies and an incomplete database for each service. The Internet has become so large that scanning even a portion of it is an extremely time-consuming process. As a result, blacklists are often not up-to-date and services are left playing catch up [3]. Furthermore, the detection algorithms used by these services may not even be completely accurate due to the development of new malicious content. Attackers continuously find new ways to hide malicious content in websites to make it undetectable by these services. Attackers have also been able to manipulate innocent websites to make them appear malicious to SBSs [4]. This causes the website to be incorrectly flagged, resulting in SBSs presenting inaccurate results to end users. Another important aspect of effective safe browsing is usability. Although users want to be protected, they still expect a smooth, non-intrusive browsing experience. The key for these services is to give users appropriate and sufficient warning without disrupting their workflow.

In light of the issues presented above, there are a number of assets at risk. Due to the extensive use of the Internet, the value of these assets is significant. Users rely on SBSs to protect sensitive information, such as banking information and other confidential data. On the other hand, commercial websites that are being inaccurately flagged risk significant damage to their reputation and website traffic, which typically results in a decrease in sales. As the exact value at risk is virtually limitless, the effectiveness of SBSs is vital.

In this report, we analyze the effectiveness of SBSs by comparing how current solutions profile potentially malicious sites. Specifically, we compare services provided by Google, Norton, and McAfee. Next, we examine how users interact with current solutions. In order to obtain practical data, we ask a group of participants to complete a survey aimed to provide us with insight into how users react to different kinds of warnings and what they prefer. Finally, from our comparison and survey results we assess a potential alternative using a prototype Google Chrome Extension.

II. Current Solutions

Safe browsing services are built into search engines (e.g. Google, Yahoo), integrated as part of web browsers (e.g. Chrome, Firefox), or offered as proprietary products (e.g. Norton Safe Web, McAfee SiteAdvisor). Most search engines show textual or graphical warnings near individual search results when a site contains suspicious content. Google takes this a step further by integrating its service with web browsers. When users visit dangerous URLs, a modal warning screen appears and forces them to either return without visiting the webpage or continue at their own risk.

Proprietary SBSs differ widely in terms of how information is displayed and how users are warned. They are generally deployed as browser plugins that show contextual icons to indicate whether the current page is safe or unsafe. A similar icon is sometimes shown beside results in selected search engines to inform users before they click a link. Some vendors use warning screens to block user from entering malicious sites, while others show warnings without preventing the user from accessing the site, and still others do nothing at all.

III. Related Work

Similar analyses of comparing the results of several SBSs have been performed by a web security company at www.stopthehacker.com. They published an article called "Website-Reputation Services Agree to Disagree" [5] in which they collected results from different SBSs using 721 suspicious websites. They found that there is a large variance between SBSs and that these services will need more work before users will be able to implicitly trust them.

IV. Analysis Methodology

A. Comparison of Current Solutions

Because there is no generic technique for detecting malicious code [6], the detection algorithms used by proprietary SBSs can produce very different results. To better understand the extent to which they differ, we developed the following technique (built upon the description here):

- Compile a list of malicious websites using publicly available sources (our "blacklist").
- Run each entry in the blacklist against selected SBSs.
- Record how the service has categorized the website.

1) Compiling the Blacklist

There are a number of public sources that maintain lists of websites suspected to contain malware, which are typically generated using a combination of proprietary software and community-based input. We chose two of these sources—Malware Patrol and Free PC Security—by virtue of their frequent updates (at least once per day) and their large list of sites (2,000 - 3,000 entries). The two lists were aggregated and filtered for duplicates to form our working, 4,952-entry "blacklist". For the purposes of this test, we assume that this blacklist is accurate and up-to-date.

2) Selecting Services to Test

To simplify development, we only used services that provided an online form or an API to access their database of malicious websites. This would allow us to quickly determine whether that service has deemed a particular website to be "safe" or "unsafe". Services that offered neither of these options—AVG, for example—would have required us to create elaborate hooks into their system, or analyze the entire blacklist by hand, both of which would have been costly to implement and execute. With this in mind, we selected three services for our test: Norton Safe Web, McAfee SiteAdvisor and Google Safe Browsing.

3) Developing and Running the Comparison

Our goal was to determine whether the selected services had marked each entry in our blacklist as "safe", "unsafe" or "unknown" (different services may have slight variations on these categories). In order to automate the procedure, we wrote Java programs to make on-demand queries for the status of all the websites from the blacklist. Upon receiving the response from the online service, the program parsed the DOM elements for specific tags, text and/or images that indicated the site's status. Results were exported to a CSV file for ease of analysis.

B. User Survey

Since current SBSs are designed to benefit users, it is necessary to understand how users interact with these services. To perform this analysis, we designed and distributed an anonymous online survey to explore how users make decisions when encountering suspicious website warnings.

1) Designing the Survey Questions

The survey focused on three topics:

1. How users proceed when encountering warnings in different situations.
2. How users decide whether to proceed or not to a potentially dangerous website.
3. Would users prefer alternative warning techniques and feel safe using them.

Specifically, users were asked how they would respond to warnings given websites of different familiarity. To explore alternative warning techniques, users were given the warning in Fig. 1, which shows an aggregated set of results for a suspicious website. The idea behind an aggregated result set is to not rely on a single service for protection but to rely on the user's ability to make an informed choice on whether to visit a suspicious website or not.

2) Distributing the Survey

The survey was distributed through email and social networking channels. These mediums allow for quick distribution and are easily accessible to our target participants. While the survey was anonymous, to understand the demographic of the data we asked participants which web browsers and SBSs they are familiar with.

Fig. 1. Aggregated warning shown to users in survey.

V. Results

A. Comparison Results

Fig. 2 shows how the different services have categorized the websites in our blacklist. Additionally:

- 6% of sites were marked as "unsafe" and 2% as "safe" by all three services.
- Of the sites Google marked as "safe", 30.6% of them are marked as "unsafe" by McAfee and/or Norton.
- Of the sites McAfee marked as "safe", 11.8% of them are marked as "unsafe" by Google and/or Norton.
- Of the sites Norton marked as "safe", 32.6% of them are marked as "unsafe" by Google and/or McAfee.

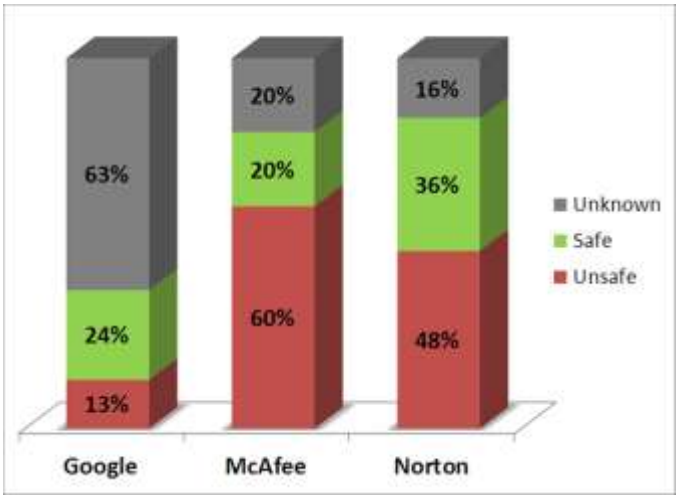

Fig. 2. Blacklist results from Safe Browsing Services.

B. Survey Results

A total of 48 anonymous users participated in the survey over the span of one week. The majority of the participants in the survey were users of the Firefox and Google Chrome web browsers. Forty five of these users have had previous encounters interacting with dangerous website warnings.

Results are shown in Tables I, II and III. (Note: results may not add up to 100% since users were allowed to skip questions and write answers other than those provided.)

When users were asked if they would proceed to a potentially dangerous website when shown the warning in Fig. 1, 81% of users said they would not proceed.

When asked if they thought it was useful to see multiple results from different services, 66% of users preferred having multiple services because they felt it either increased the credibility of the warning and that more information was useful to them. Thirty four percent of users did not find having multiple services useful because they found it confusing or would rather have a verdict on a website rather than a choice.

When asked if they thought it was useful to see when the last scan was performed on a website, 87% of users found it useful because they believe that more recent results would be more accurate. Thirteen percent of users were confused by the last scan date information.

Users were given a scenario in which a website was not found to be malicious, but contained links to potentially dangerous websites. The users were asked if they would prefer to (1) see a warning message before entering the website or (2) have their browser clearly identify the dangerous links to the user while viewing the website.

TABLE I
HOW USERS PROCEED WHEN ENCOUNTERING A SUSPICIOUS WEBSITE WARNING

Action	Unfamiliar Website	Familiar Website
Do not proceed	63%	4%
Need to verify before proceeding	33%	65%
Proceed to website	4%	31%

TABLE II
HOW USERS' IMPRESSIONS OF A WEBSITE IS AFFECTED BY WARNINGS

Action	Unfamiliar Website	Familiar Website
Do not trust website; cautious about ever visiting again	58%	6%
Trust website only if no warning is displayed	29%	60%
Trust website even if warning is displayed	0%	17%

TABLE III
USERS' PREFERENCES FOR WARNINGS

(1) Warning before entering a website	27%
(2) Warnings while viewing a website	33%
Both (1) and (2)	38%

VI. Discussion

A. Comparison Discussion

Guaranteed Diversity - There is clearly some diversity in how each service has profiled the blacklist. As mentioned, services employ different methods of detecting malicious code. They may scan entire pages or only parts of a site; they can detect suspicious behaviors as they occur or merely look for their "postmortem" effects (e.g. virus signatures). Additionally, services may not have the most updated status of a website.

Safety In Numbers - McAfee detected the largest portion of unsafe sites (60%), but it does not necessarily indicate that it is more accurate than the rest. Nearly 12% of what it marked as "safe" was marked as "unsafe" by Norton and/or Google. In fact, in 4,952 sites, the services only reached a consensus on 439 (8%) of them. With such a large discrepancy between the results, any single service will offer only limited protection against possible threats.

Limited Coverage - The high number of "unknown" results is also a concern, as they represent a significant vulnerability for users who may assume a website is safe as long as a warning does not appear. It is also an indication of the services' coverage; Google, for instance, does not seem to scan nearly as many websites as Norton and McAfee. While it is difficult to determine each service's exact coverage, it is clear that their lists are neither complete nor perfectly overlapping.

Data Aggregation - It is possible to overcome the limited accuracy and coverage of SBSs by aggregating their results. For instance, whereas Google, McAfee and Norton only have results for 37%, 80% and 84% of the blacklist, respectively, a list that combines the three will cover 97% of the blacklist. At the same time, the results for a particular site are also verified by multiple sources before it is deemed safe or unsafe.

B. User Survey Discussion

Limitations of Results - The user study was designed to be a preliminary user study and not designed to be statistically meaningful. As such, results from the study should be used to perform a larger more focused study.

Flaws In Warnings - Our survey shows that when websites are flagged as suspicious by SBSs, the consequences for website owners are non-trivial. Tables I and II show that when users encounter warnings for unfamiliar websites, the majority will automatically consider it dangerous and possibly never visit again. While the results for users navigating to familiar websites are not as ominous, the fact that users are less likely to trust warning messages on familiar websites leave users vulnerable to actual attacks.

Intelligent Decisions - Users were shown the warning in Fig. 1, which doesn't directly label a website as safe or unsafe. Instead the users were given aggregated information from different sources to aide users in deciding whether or not to proceed to a website. The results showed that users were generally perceptive to the added information, but were not necessarily able to apply that information effectively. While the results do not state that giving users more information will necessarily provide a safer and more pleasant online experience, it does suggest that the majority of users could apply information intelligently where SBSs currently do not.

C. Additional Observations

Single Choice – Another interesting choice that all SBSs employ is providing only a single, binary choice for users to make. This relates to an important security principle—Defense in Depth. As soon as a user selects to continue to a website, SBSs disengage and allow the user do as they wish. Instead, SBSs should provide a second or third level of defense by blocking out links or sandboxing the website even if the user chooses to continue.

Psychological Acceptability – Another aspect or security principle of concern with current SBSs relates to how they warn users of malicious websites. When users receive a site warning, they are given a binary choice of whether to continue to the website or return to the previous safe site without much insight on the particular situation other than the fact that the website "may be malicious." This can be confusing to a user who may have been visiting the site regularly without issues. If SBSs provided further information, such as highlighting specific links or content on the website that is unsafe, users may have a better understanding of the situation and will be able to make a more informed decision.

D. Possible Countermeasures

1) Design Principles

Based on our analysis of SBSs offered by Norton, McAfee and Google, it is clear that these services had missing components or design principles that we would have liked to see implemented. These are: (1) aggregated data regarding malicious sites and (2) a simple, non-intrusive approach to identify malicious links on a website. While the focus of this project is to provide an in-depth analysis on current solutions, we decided that it would be worthwhile to develop a small prototype application to test our findings.

We decided the ideal interface for this solution would be as a plug-in for an existing browser —specifically Google Chrome. Doing so allowed us to easily integrate our tool with the browser and directly interact with website elements. LinkGuard, the Chrome extension that we developed, is written in JavaScript and utilizes the chrome.extension module.

2) Data Source and Aggregation

The first goal for our prototype is to aggregate malicious site data from as many reliable sources as possible, resulting in a more comprehensive database of websites and therefore strengthening fail-safe defaults for the user. There are two possible approaches to verify whether a site is malicious or not: active and passive.

Using an active approach, the plugin would check the status of each link on a page "on-the-fly." In other words, for each

link on a website, the plugin would send a request to all data sources and display the results on each and every page load. Although this would be ideal for retrieving the most up-to-date result, we concluded that an active approach would consume an unreasonable amount of bandwidth and processing time for each page load, and decided against it for this first prototype. Regardless, we believe a more active approach should be reconsidered at a later time if a more efficient scanning and parsing process could be implemented.

In a passive approach, the plugin would check the status of a predetermined list of malicious sites at a predefined interval. Ideally, at each interval the plugin would update its list of malicious sites and the status for each. In our approach, we decided to use the same list of malicious sites that we had previously compiled in our analysis stage. As McAfee, Norton and Google are well-known, reputable companies in the Internet Security space, we felt this list was reasonable for our initial prototype. Due to time constraints, this list of malicious websites was compiled and parsed only once.

3) Non-Intrusive Notification of Malicious Links

The second goal for our prototype application is to identify malicious links and display a warning to users in a non-intrusive manner. As our results suggest, modal warning screens used by current SBSs are only beneficial if the entire site is malicious. However, for legitimate websites that have been accidentally flagged as malicious due to a few user-posted links, this approach can be devastating to the site's reputation. Based on result from our survey, we wanted to develop a solution that would protect users from malicious links, but at the same time, would not interrupt their browsing experience.

The prototype plugin warns users of malicious links in two forms: a non-modal browser message (Fig. 3) and an in-line warning of malicious links (Fig. 5). The non-modal browser message is used instead of a warning splash page because it is less disruptive to the user but still provides a sufficient warning of a potentially dangerous page. Additional information on which links are potentially malicious can be viewed by clicking on the plugin icon (Fig. 4). As mentioned, another limitation that other SBSs share is that they warn the user about the page only once. After the user has entered the page, these services do not specifically inform the user what to watch out for. By automatically disabling malicious links and providing an in-line warning to the user, the plugin takes a non-intrusive approach and gives the user much more context and information about the problem.

Malicious links are overlapped with red strikes to draw the user's attention. When the user hovers the mouse pointer over the crossed out link, the red strikes disappear and the original URL becomes visible. If the user decides to click on the provided link, the system displays a prompt as a final warning (Fig. 6). This multi-layered approach provides defense in depth, and forces users to confirm their choices before they become exposed to malicious content.

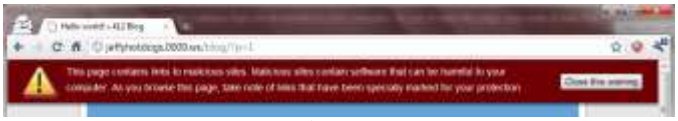

Fig. 3. Non-modal browser notification popup.

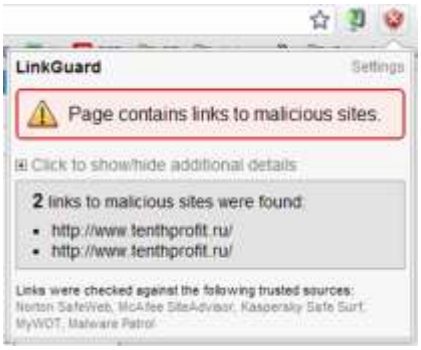

Fig. 4. LinkGuard plugin menu.

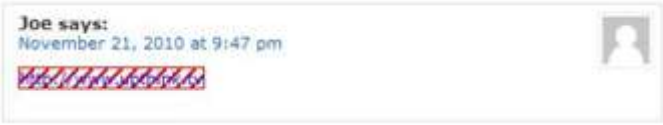

Fig. 5. LinkGuard automatically disables links.

Fig. 6. LinkGuard final warning after clicking malicious link.

4) Reception and Summary

We received preliminary feedback from numerous friends and family members to whom we introduced LinkGuard, and all responses were very positive. Many of those surveyed would consider using the plugin regularly if it were further refined and polished.

In summary, LinkGuard achieves the two main goals we had in mind. Although there are definitely enhancements and optimizations that can be made, LinkGuard is a successful prototype that sufficiently demonstrates a refined approach to SBSs. By aggregating multiple data sources, LinkGuard is able to warn users about significantly more potentially unsafe websites than any single service. Also, by providing a non-intrusive, in-line warning of malicious links, LinkGuard improves the end user experience while still maintaining the same level of security as before.

E. Ideal Countermeasure

In our analysis, we identified numerous techniques and methods that we would have liked to see in an ideal SBS. In addition to what was demonstrated in LinkGuard, the

following techniques should be considered in an ideal countermeasure:

1. Immediate "on-the-fly" scanning of websites against all data sources.
2. Combination of professional data sources (Norton, McAfee, Google) and community-based data sources (user-submitted reviews of website).
3. "Sandboxing" of malicious sites for better security and to prevent drive-by malware downloads [7].
4. Full browser integration without having to download and enable a plugin.

VII. Conclusion

Based upon our findings and analysis, it became clear to us that the current solutions to ensure safe browsing are inadequate and lack important aspects that can improve security and users' browsing experience. By compiling a list of potentially malicious websites and testing it against three prominent SBSs, we concluded that relying on a single service to protect a user from malicious attacks is insufficient. Instead, an aggregated approach offers a significant advantage in identifying potentially malicious links. In order to obtain practical and current data, we also conducted a survey to help us better understand how users interact with the current solutions. From the results of a diverse group of participants, it is evident that intrusive warnings as offered by the current solutions leave a negative impression of the website to end users. Participants also prefer to see detailed information from a variety of different services to help them make a better judgment on whether a site is indeed malicious. Finally, we demonstrated the viability of a SBS that employed aggregated results and less intrusive warnings using a prototype.

References

[1] Vamosi, R. (2010, Aug.). "Cross-Site Scripting: An Old Problem Returns." *PC World*. [Online]. 28(8), pp. 37-38. Available: web.ebscohost.com [Dec. 6, 2010].

[2] Dyrli, K. (2009, Mar.). "Security Software: Warding Off Viruses." *District Administration*. [Online]. 45(3), pp. 42. Available: web.ebscohost.com [Dec. 6, 2010].

[3] Chen, T and Wang, V. (2010, Mar.). "Web Filtering and Censoring." *Computer*. [Online]. 43(3), pp. 94-97. Available: web.ebscohost.com [Dec. 6, 2010].

[4] Krebs, B. "Hiding from Anti-Malware Search Bots." Internet: http://krebsonsecurity.com/2010/04/hiding-from-anti-malware-search-bots/, Apr. 23, 2010 [Dec. 6, 2010].

[5] "Website-Reputation Services Agree to Disagree" Internet: http://www.stopthehacker.com/2010/01/17/website-reputation-services-agree-to-disagree/, Jan. 17, 2010 [Dec. 6, 2010].

[6] Beznosov, K. Eece 412. Class Lecture, Topic: "Malicious Software." MCLD 254, Faculty of Engineering, University of British Columbia, Vancouver, British Columbia, Nov. 2, 2010.

[7] Wallach, D. (2010, Jan.). "Native Client: A Clever Alternative." *Communications of the ACM*. [Online]. 53(1), pp. 90. Available: web.ebscohost.com [Dec. 6, 2010].