# Do Security Toolbars Actually Prevent Phishing Attacks?

Min Wu, Robert Miller, Simson Garfinkel
MIT CSAIL
32 Vassar Street
Cambridge, MA 02139
{minwu, rcm, simsong}@csail.mit.edu

## ABSTRACT

Security toolbars inside a web browser are designed to show security-related information about a website in order to help users detect phishing websites. We ran a user study to test the effectiveness three security toolbars preventing phishing attacks. The results show that all the tested security toolbars are ineffective: users were spoofed 34% of the time. Even users were asked to pay attention to the toolbar, if the content of web pages is good enough, some users will disregard the toolbar display. Moreover, since many companies do not follow good practice in designing their websites, the toolbar cannot help some users to distinguish poorly designed websites from malicious phishing attacks.

## 1. INTRODUCTION

Phishing uses emails and websites, designed to look like they come from legitimate organizations, to deceive users into disclosing their personal or financial information. Phishing is rapidly increasing and seems to be successfully fooling users. According to the Anti-Phishing Working Group [2], the number of active phishing sites reported in March 2005 is 2870. The average monthly growth rate in phishing sites from July 2004 to March 2005 is 28%. A recent study [3] done by the anti-spam firm MailFrontier Inc. found that phishing emails fooled users 28% of the time.

Many implemented anti-phishing schemes use a *security toolbar*, which is a toolbar inside the web browser that displays security information about the website. SpoofStick [4] displays the website's domain name. Netcraft Toolbar [7] displays more site information including the site's registration date, popularity, hosting country, and the name of the netblock on which the site is hosted. Trustbar [1] makes SSL connections obvious by displaying logos of the website and its Certificate Authority based on the principle that legitimate websites use SSL to encrypt the user's sensitive data transmission, while phishing sites do not bother to use SSL. eBay's Account Guard [5] uses a red, green or gray icon to indicate whether a site is known phishing, truly belonging to eBay or PayPal, or belonging to others. SpoofGuard [6] calculates a spoof score for the current web page based on a set of heuristic detection rules derived from the common characteristics of known phishing attacks. It then displays a red, green, or yellow light to indicate that the page is analyzed to be phishing, good, or unsure.

Security tips on phishing attacks ask users to always pay attention to these security toolbars whenever they access a web site. But are these toolbars actually effective at preventing phishing attacks? We performed a user study in order to test these security toolbars and, more generally, to find out why users get fooled by phishing attacks.

## 2. STUDY DESIGN

Instead of testing individual real toolbar, we designed three abstract security toolbars based on the existing ones (see Figure 1) because there are three types of information displayed by the existing toolbars. The *Neutral-Information toolbar* integrates SpoofStick and Netcraft Toolbar and displays neutral administrative information about the current website including the domain name, the domain's registration date, and the domain's hosting country. The *SSL-Verification toolbar* imitates Trustbar and displays the confirmation information for secure site but the warning message for other sites. The *System-Decision toolbar* simulates eBay Account Guard and SpoofGuard, presenting a judgment about the site's trustworthiness using a red light, a green light or a yellow light to indicate that the sites are analyzed to be trustworthy, phishing or unsure, respectively. When the toolbar displays the red light, it also shows a red text message of "Potential Fraudulent Site" as an explanation.

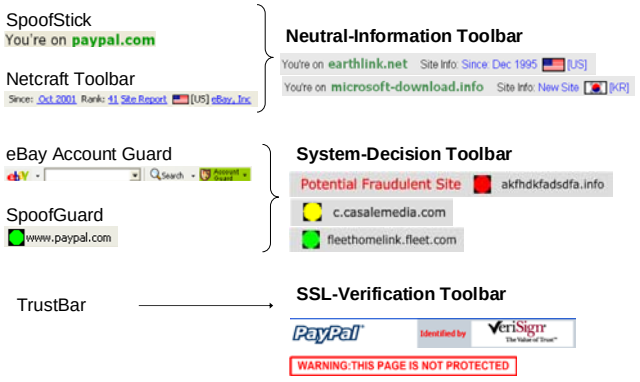


**Figure 1. Three abstract toolbars**

We first set up dummy accounts as John Smith at various websites. Users were told that they are John's personal assistant and their tasks were to manage John's wish lists at those sites which are protected by his username and password. Users dealt with 20 emails sent from various websites and then forwarded by John Smith, five of which were phishing emails. Users had to click the link in the forwarded emails in order to go to the corresponding websites.

In the study, this link always led to the *real* website, regardless of whether the email was phishing or not. To simulate phishing attacks, we changed the appearance of the security toolbar and other browser security indicators like the address bar to indicate

that the web page was served by an unusual source, *e.g.*, amazon-department.com rather than amazon.com. As a result, our study simulated *ideal* phishing attacks whose content is a perfect copy of the actual website. In practice, an attacker might not bother mirroring the entire site, but simply act as a man-in-the-middle between the user and the actual site.

There were five types of phishing attacks. Each user saw one of each. The similar-name attack, IP-address attack, and hijacked-server attack replaced the real site's hostname with a similar hostname, an IP address or a totally different hostname, respectively. The popup-window attack used a popup window for user's login information while the real site is displayed in the background. The Paypal attack copied a real phishing attack on Paypal.

All three toolbars were configured to discriminate between the legitimate sites and the phishing sites. For example, none of the phishing sites used SSL so the SSL-Verification toolbar always displayed a warning for them. For the System-Decision toolbar, all legitimate sites were trustworthy but all the phishing sites were displayed as phishing or unsure.

There was a tutorial email halfway through the study. The email introduced to the users what a phishing attack is and told them how to use the tested security toolbar to detect phishing attack.

A total of 30 users were recruited at MIT, with 14 females and 16 males. Twenty of them were MIT students from 10 different majors. The average age of the users was 27 with the range from 18 to 50. Ten users were assigned to each security toolbar.

## 3. RESULTS AND DISCUSSION

The *spoof rate* is the percentage of the simulated attacks that successfully got John's username and password or other personal information without raising the user's suspicion. All the tested toolbars have high spoof rates (see Figure 2). The average spoof rate in the study is 34%.

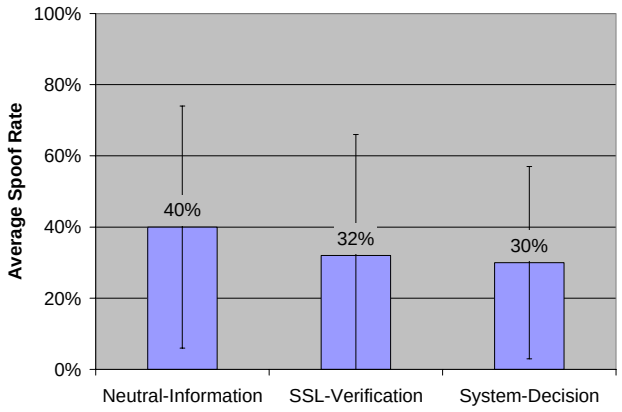


**Figure 2. Spoof rates of different toolbars**

The difference between the spoof rates before and after users seeing the tutorial is marginally statistically significant ($p$ = 0.063). The spoof rate before the tutorial is 43%, while the spoof rate after the tutorial is 25%. The decrease of the spoof rates before and after the tutorial is also maintained among all three toolbars. Several users mentioned that the tutorial email helps them to pay attention to the security toolbars. Another contribution to this result is that the spoof rate steadily decreases as users experience more and more similar attacks.

The attack type significantly affects the spoof rate ($p$ = 0.052). The Paypal attack has the lowest spoof rate (17%) because it asked for sensitive financial data and it imitated real phishing attacks that some users have already seen. The popup-window attack has the second lowest spoof rate (28%) because a majority of users tend to close the popup window without even reading it. The other three attacks (similar-name, IP-address, and hijacked-server) have high spoof rates of 50%, 33%, and 43%, respectively.

If the Paypal attack is not considered, the attack position significantly affects the spoof rates ($p$ = 0.039). The later the attack happens, the lower the spoof rate is, from 57% at the 1st attack to 23% at the 5th attack. This result shows that users do learn how to detect similar phishing attacks as they experience more of them.

Twenty out of 30 users got spoofed by at least one phishing attack. Among the 20 users, 17 (85%) mentioned that the reason they were fooled is that the web content looks professional or exactly the same as they visited before. This indicates that even though the security toolbars express suspicious signs and users do notice them, if the web content is good enough, some users disregard those suspicious signs. Moreover, 8 users (40%) explained away phishing URLs with reasons like "it may be an outsourcing site" or "sometimes I go to a website and the site directs me to another address which is different from the one I typed in." This indicates that poorly designed websites make phishing attacks worse. Some legitimate companies do not have consistent domain names for their web sites. Some register vague or unrelated domain names. Outsourcing is also a big problem. As a result, users cannot distinguish poorly designed websites from the malicious phishing attacks.

## 4. REFERENCES


[1] Amir Herzberg and Ahmad Gbara. *TrustBar: Protecting (even Naïve) Web Users from Spoofing and Phishing Attacks.* http://www.cs.biu.ac.il/~herzbea/Papers/ecommerce/spoofing.htm

[2] Anti-Phishing Working Group. *Phishing Activity Trends Report, March 2005.* http://antiphishing.org/APWG_Phishing_Activity_Report_March_2005.pdf

[3] Bob Sullivan. *Consumers still falling for phish.* MSNBC, July 28, 2004. http://www.msnbc.msn.com/id/5519990/

[4] CoreStreet. *SpoofStick.* http://www.corestreet.com/spoofstick/

[5] eBay. *eBay Toolbar and Account Guard.* http://pages.ebay.com/help/confidence/account-guard.html

[6] Neil Chou, Robert Ledesma, Yuka Teraguchi, and John C. Mitchell. *Client-Side Defense Against Web-Based Identity Theft.* 11th Annual Network and Distributed System Security Symposium, 2004. http://theory.stanford.edu/people/jcm/papers/spoofguard-ndss.pdf

[7] Netcraft. *Netcraft Toolbar*. http://toolbar.netcraft.com/