

User Behaviour Based Phishing Websites Detection

Xun Dong

Department of Computer Science
University of York
York, United Kingdom
Email: xundong@cs.york.ac.uk

John A. Clark

Department of Computer Science
University of York
York, United Kingdom
Email: jac@cs.york.ac.uk

Jeremy L. Jacob

Department of Computer Science
University of York
York, United Kingdom
Email: jeremy@cs.york.ac.uk

Abstract—Phishing detection systems are principally based on the analysis of data moving from phishers to victims. In this paper we describe a novel approach to detect phishing websites based on analysis of users' online behaviours – i.e., the websites users have visited, and the data users have submitted to those websites. Such user behaviours can not be manipulated freely by attackers; detection based on those data can not only achieve high accuracy, but also is fundamentally resilient against changing deception methods.

Index Terms—Phishing Attacks, Phishing Websites Detection, Identity Theft, User Protection

I. INTRODUCTION

PHISHING attacks are well-organised and financially motivated crimes which steal users' confidential information and authentication credentials. They not only cause significant financial damage to both individuals and companies/financial organisations, but also damage users' confidence in e-commerce as a whole. According to Gartner analysts, financial losses stemming from phishing attacks have risen to more than 3.2 billion USD with 3.6 million victims in 2007 in US [23], and consumer anxiety about Internet security resulted in a two billion USD loss in e-commerce and banking transactions in 2006 [22].

The scale and sophistication of phishing attacks have been increasing steadily despite numerous countermeasure efforts. The number of reported phishing web sites increased five-fold from 10047 to 55643 in the 10 month period between June 2006 and April 2007 [3]. The real figure may be much higher because many sophisticated phishing attacks (such as context aware phishing attacks, malware based phishing attacks, and real-time man-in-the-middle phishing attacks against one-time passwords [34]) may not all have been captured and reported.

In this paper we present our design, implementation and evaluation of the *user-behaviour* based phishing website detection system (UBPD). UBPD does not aim to replace existing anti-phishing solutions, rather it complements them. It alerts users when they are about to submit credential information to phishing websites (i.e. when other existing countermeasures fail), and protects users as the last line of defence. Its detection algorithm is independent from how phishing attacks are implemented, and it can easily detect sophisticated phishing websites that other techniques find hard to deal with. In contrast to existing detection techniques based only on the incoming data, this technique is much simpler, needs to deal with much

less low level technical detail, and is more difficult to bypass by varying spoofing techniques and deception methods.

Note: In the rest of the paper we use 'interact', 'interaction' and 'user-webpage interaction' to refer to the user supplying data to a webpage.

II. RELATED STUDIES

Researchers have studied the characteristics of phishing attacks [29], [30], [38] and have provided models of phishing attacks [4], [10], [18]. Our understanding of human factors in phishing has been improved by works such as [9], [17], [19], [20], [33]. Novel security interfaces, and automated detection systems have also been invented to protect users.

Anti-phishing email filters can fight phishing at the email level, as it is the primary channel for phishers to reach victims. SpamAssassin [2], PILFER [11], and Spamato [5] are typical examples of those systems. They analyse the incoming emails by applying predefined rules and characteristics often found in phishing emails. PHONEY is a phishing email detection system that tries to detect phishing emails by mimicking user responses and providing fake information to suspicious web sites that request critical information. The web sites' responses are forwarded to the decision engine for further analysis [7].

Web Wallet [40] creates a unified interface for authentication. Once it detects a login form, it asks the user to explicitly indicate the intended site to login. If the intention matches the current site, it automatically fills webpage input fields. Otherwise a warning will be presented to the user.

Phishing website filters in Internet Explorer [24], safe browsing in Firefox 2/3 [27], and Netcraft toolbar [28] are all blacklist anti-phishing websites detection systems. They check whether the URL of the current web page matches any identified phishing web sites before rendering the webpage to users.

SpoofGuard [8] is a signature-and-rule based detection system. It analyses the host name, URL, and the images used in the current webpage to detect the phishing websites. CANTINA [41] uses the TF-IDF information retrieval algorithm to retrieve the key words of the current webpage, and uses Google search results to decide whether the current website is phishy.

Ying Pan et al. have invented a phishing website detection system which examines the anomalies in web pages, in par-

ticular, the discrepancy between a web site's identity and its structural features and HTTP transactions [31].

III. DESIGN

A. Phishing Nature and Detection Philosophy

Phishing websites work by impersonating legitimate websites, and they have a very short life time. On average a phishing website lasts 62 hours, and users rarely visit the phishing website prior to the attack [26].

Secondly, phishing attacks always generate mismatches between a user's perception and the truth. In successful web based phishing attacks, victims have believed they are interacting with websites which belong to legitimate and reputable organisations or individuals. Thus the crucial mismatch that phishers create is one of real *versus* apparent identity.

Phishing attacks can be detected if we can detect such a mismatch. One approach is to predict a user's perception and then compare it with the actual fact understood by the system. CANTINA is an example of this approach [41]. The main issue with this approach is that the data the system relies on is under the control of attackers, and there are so many techniques that attackers can apply to manipulate the data to easily evade the detection. For CANTINA, attackers could use images instead of text in the body of the webpage, they could use iframes to hide a large amount of content from the users while computer programs can still see it; they could use Javascript to change the content of the page after the detection has been done.

We decide to use another, more reliable, approach. The authentication credentials, which phishers try to elicit, ought to be shared only between users and legitimate organisations. Such (authentication credential, legitimate website) pairs are viewed as the user's *binding relationships*. In legitimate web authentication interactions, the authentication credentials are sent to the website they have been bound to. In a phishing attack the mismatches cause the user to unintentionally break binding relationships by sending credentials to a phishing website. No matter what spoofing techniques or deception methods used, nor how phishing webpages are implemented, the mismatch and violation of the binding relationships always exists. So one can discover the mismatch by detecting violation of users' binding relationships.

Hence phishing websites can be detected when both of the following two conditions are met: 1) the current website has rarely or never been visited before by the user; 2) the data, which the user is about to submit, is bound to website other than the current one.

B. System Design

1) *Overview of the detection work flow:* UBPD has three components:

- **The user profile** contains data to describe the users' binding relationships and the users' personal whitelist. The profile must be constructed before the system can detect phishing websites.

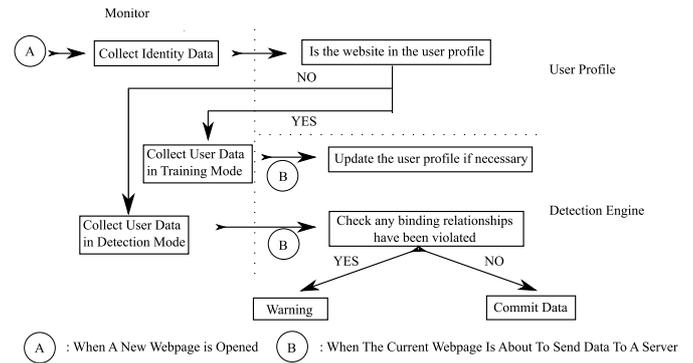


Fig. 1. Detection Process Work Flow

- **The monitor** collects the data the user intend to submit and the identity of the destination websites, and activates the detection engine.
- **The detection engine** uses the data provided by the monitor to detect phishing websites and update the user profile if necessary.

UBPD has two working modes: training mode and detection mode. In training mode, UBPD runs quietly in the background, and focuses on learning newly created binding relationships or updating the existing binding relationships. Only in detection mode, UBPD checks whether any of the user's binding relationships would be violated if the user submitted data is sent to the current website. The mode in which UBPD runs is decided by checking whether the webpage belongs to a website

- 1) whose top level domain¹ is in the user's personal whitelist, or
- 2) with which the user has shared authentication credentials.

If either is true the system will operate in training mode, otherwise, in detection mode. Phishing webpages will always cause UBPD to run in the detection mode, since they satisfy neither condition.

The detection work flow is shown in Figure 1. Once a user opens a new webpage, the monitor decides in which mode UBPD should be running. Then, according to the working mode the monitor chooses appropriate method to collect the data the user submitted to the current webpage, and sends it to the detection engine once the user initiates data submission. The details of the data collection methods are discussed in section IV. When running in detection mode if the binding relationships are found to be violated, the data the user submitted will not be sent and a warning dialogue will be presented. For the remaining cases, UBPD will allow the data submission.

2) *Creation of the User Profile:* The user profile contains the user's binding relationships and personal whitelist. The binding relationships are represented as a collection of paired records, i.e., $\langle aTopLevelDomain, aSecretDataItem \rangle$. The personal whitelist is a list of top level domains of websites.

¹Suppose the URL of a webpage is "domain2.domain1.com/files/page1.htm", the top level domain is "domain1.com"

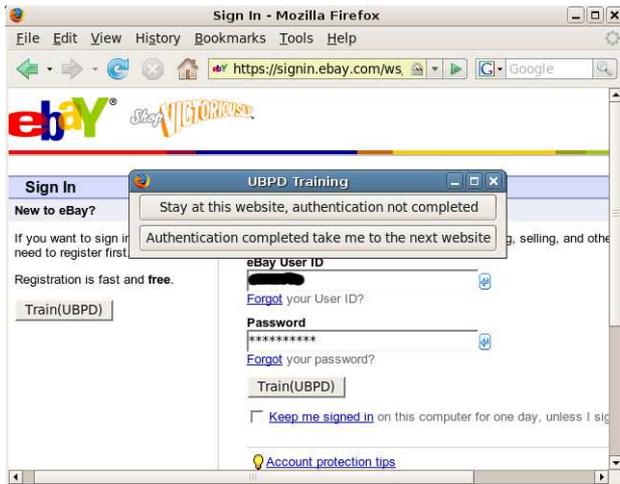


Fig. 2. User Profile Creation

To make sure the binding relationship records are created, UBPD makes it as a mandatory step of UBPD's installation procedure. Having installed the UBPD, when the browser is restarted for the first time, the user is presented with a dialogue, which asks for the websites with which the user has accounts. Then UBPD automatically takes users to those websites, and asks users to fill in the authentication forms on those websites one by one. In addition, UBPD also dynamically modifies each authentication webpages, so that all the buttons on the webpage will be replaced with the 'Train(UBPD)' button. Once a user clicks on it, UBPD creates new binding relationship records in the user profile and then asks whether the user wants to carry on the authentication process (in case it is a multiple step authentication) or go to the next website. A screen shot is shown in Figure 2.

On average users have over 20 web accounts [12]. We do not expect them to train UBPD with all their binding relationships, however, we do expect them to train UBPD with their most valuable binding relationships (such as their online accounts with financial organisations). Our preliminary evaluations have confirmed that our assumption is valid.

There are two types of binding relationships that UBPD is unaware of: the ones that already exist but users have not trained UBPD with, and the ones that users will create in the future. It is possible that the authentication credentials in these unknown binding relationships have also been used in the binding relationships that UBPD is aware of (as the result of authentication credentials reuse). When this is the case, with the binding relationships records alone UBPD interprets the legitimate use of the unknown binding relationships as a violation of binding relationship, and wrongly raise a phishing attack warning. To avoid such false warnings, the whitelist is automatically created to include the websites that users have already got accounts with but did not make UBPD aware of and the websites they may have account with in future. As stated in section III-B1 if the website with which the user is currently interacting is found to be in the whitelist UBPD

will only run in training mode. In the training mode, such reuse of the authentication credentials can only cause UBPD to either create new binding relationships record or update existing ones.

The personal whitelist is constructed by combining a default whitelist with the websites the user has visited more than three times (configurable) according to the user's browsing history. The default whitelist is constructed by identifying the 1000 websites that are most visited in the user's country. This information is obtained from Alexa, a company specialising in providing internet traffic ranking information. The default whitelist could be viewed as a reflection of mass users' online browsing behaviours. These most popular websites are not phishing websites, and users are likely to have binding relationships with them in the future if they do not have already. In this way, websites that users currently have accounts with or are going to have account with are very likely to be included in the whitelist. By default, this whitelist is automatically updated weekly. It is worthy noting that, the whitelist in UBPD is mainly used for reducing false warnings, it is not used as conventional whitelist to decide whether the current website is legitimate or not.

3) *Update of the User Profile:* Besides the manual update by users, when running in the training mode UBPD has an automatic method to update the user profile with the other unknown binding relationships. It detects whether the user is using a web authentication form. This is achieved by analysing the HTML source code, such as the annotation, label, use of certain tags (such as `<form>`) and type of the HTML elements. If the user is using a web authentication form, and the user profile contains no binding relationships with the current website, then UBPD prompts a window to ask the user to update the user profile. If there is an existing binding relationship for the current website, then UBPD will replace the authentication credentials in the binding relationships with the latest values the user submits. If users have entered the authentication credentials wrongly, those credentials will still be stored, but those wrong values will be corrected, when users relog in with the correct authentication credentials. In future the detection of web authentication page usage can be much simpler and more accurate once web authentication interfaces [15], [35], [37] are standardised. Current authentication form detection is still accurate because it deals with legitimate websites (automated update is carried out only in training mode), and those legitimate websites use standard web forms and will not deliberately disrupt detection.

4) *Phishing Score Calculation:* In detection mode, UBPD decides whether the current webpage is a phishing webpage by calculating phishing scores. The calculation is a two step process. In the first step, for each legitimate website, with which the user has shared authentication credentials, a temporary phishing score is calculated. Each temporary phishing score is the fraction of the authentication credentials associated with a legitimate website that also appear in the data to be submitted to the current webpage. Its value ranges from 0.0 to 1.0.

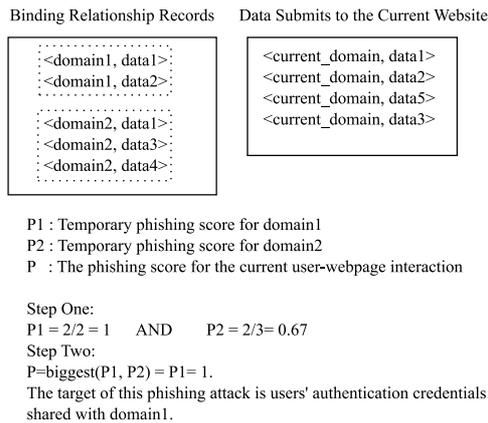


Fig. 3. An Example of How Phishing Score Is Calculated

In the second step, those temporary phishing scores are sorted into descending order. The current webpage's phishing score is the highest score calculated. The legitimate website with the highest temporary phishing score is considered to be the impersonated target of the phishing website. If more than one legitimate website has yielded the highest phishing score (due to the reuse of authentication credentials), they will all be considered as targets. Although it may not be the attacker's intent, the data they get if an attack succeeds can certainly be used to compromise the user's accounts with those legitimate websites. Figure 3 illustrates how a phishing score is calculated in UBPD.

Given our phishing score calculation method, clever attackers may ask the user victims to submit their credential information through a series of webpages, with each phishing webpage asking only for a small part of data stored in the user profile. To handle this fragmentation attack UBPD has a threshold value and cache mechanism. Once the phishing score is above the threshold the current webpage will be considered as a phishing webpage. The system's default threshold is 0.6. Why we choose 0.6 is discussed in section V. If the current phishing score is not zero, UBPD also remembers which data has been sent to the current website and will consider it in the next interaction if there is one. Accordingly many fragmentation attacks can be detected.

5) *Reuse*: It is very common that a user shares the same authentication credentials (user names, passwords, etc) with more than one website. The two running modes and the user's personal whitelist are designed to prevent false warnings caused by reuse without compromising detection accuracy.

UBPD detects the violation of binding relationships only when the user is interacting with websites that are neither in the user's whitelists nor which the user has account with. So as long as legitimate websites for which users have used the same authentication credentials are all contained in the user profile, there will be no false phishing warning generated due to the reuse. The method that UBPD uses to create the user profile ensures such legitimate websites are most likely to be included, as those websites are either within the user's

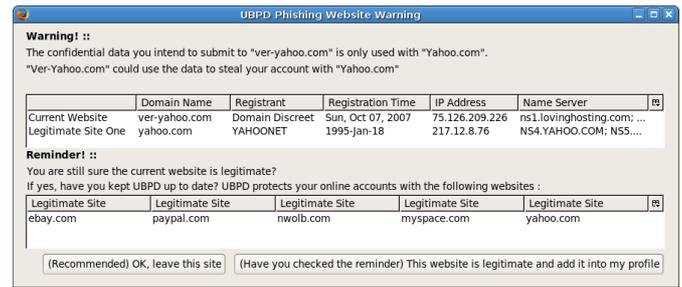


Fig. 4. Phishing Warning Dialogue

browsing history or are popular websites in the user's region. Our preliminary false positive evaluation also confirmed this.

6) *Warning Dialogue*: The content of the warning dialog must be suitable for users with very limited knowledge, otherwise, as previous user study [39] found, users may ignore the warning or may not behave as suggested. Figure 4 is a warning dialog example. To make the information easy to understand, the dialogue tells users that the current website, to which they are submitting credentials, is not one of the legitimate websites associated with those authentication credentials. To help users understand the detection result and make a correct decision, UBPD also provides information regarding the differences between the legitimate website and the possible phishing website in five aspects: the domain name, the domain registrant, the domain registration time, name servers, and IP addresses. Users don't need to understand those terms. They only need be able to recognise the difference between values of the five attributes of the legitimate website and the phishing website.

7) *Website Equivalence*: To discover whether the user is about to submit the authentication credentials to entities, with which they have not been bound, UBPD needs to be able to accurately decide whether a given website is equivalent to the website recorded in user's binding relationships. It is much more complicated than just literally comparing two URLs or IP addresses of two websites, because: 1) big organisations often have web sites under different domain names, and users can access their account from any of these domains; 2) the IP address of a website can be different each time if dynamic IP addressing is used; 3) it is hard to avoid 'pharming' attacks, in which the phishing site's URL is identical to legitimate one.

Our system first compares two websites' domain names and IP addresses. When the two domain names and two IP addresses are equal the web sites are assumed to be identical. Otherwise, the system interrogates the WHOIS² database and uses the information returned to determine equivalence. When analysing two different IP addresses our system compares the netnames, name servers, and the countries where each IP address is registered. If they are all identical then the two websites are deemed to be identical. This method can

²WHOIS is a TCP-based query/response protocol which is widely used for querying a database in order to determine the owner of a domain name, an IP address. RFC 3912 describes the protocol in detail.

also detect pharming attacks, in which both fraudulent and legitimate websites have the same domain name but are hosted on different IP addresses.

This method is not perfect. A more elegant and complete solution would be a mechanism where servers provide security relevant metadata (including the outsourcing information) to the web browser via a standard protocol as suggested by Behera and Agarwal [6]. However, unless it becomes a standard we have to rely on WHOIS. The extended validation certificate [13] can provide information to decide whether two websites belong to the same party, but the problem is the high cost, high entry level and complexity of obtaining such certificates. Many small and medium businesses will not be able to own them; as a result they cannot be used to decide website equivalence in general.

8) *User's Privacy*: Since a user profile contains confidential user information, it is important that it does not add new security risks. We use a one-way secure hash function to hash the confidential data before it is stored. When the system needs to determine the equivalence between the data, the system just needs to compare the hash values.

However, because the domain name of the website is not hashed, if the profile is obtained by attackers they would be able to find out with which websites users have accounts and where users have reused their authentication credentials. This information is helpful for attackers; for example, it can be used to launch context aware phishing attacks against the users. To prevent this, when the system is installed it will randomly generate a secret key, and this key will be used to encrypt and decrypt the domain names.

C. Implementation

UBPD is implemented as a Firefox add-ons. Most of the detection is implemented using Javascript. The hash function we use is SHA-1[21] and the encryption method we use is Twofish[1]. The hash function and encryption methods can be changed, we choose to use them mainly because there are open source implementations available. The user interface of the system is implemented by using XUL, a technology developed by Mozilla [14]. Although only implemented on Firefox, the detection system can be easily developed to work on other browsers, such as Internet Explorer and Opera.

IV. EVASION AND COUNTERMEASURES

A. Disruption of Data Collection

To accurately detect phishing UBPD needs to be able to work with the exact data the user is submitting to the current webpage. This data can be easily retrieved by accessing the DOM interface, however, the client script language can also manipulate the data before the monitor retrieves it.

In training mode, websites that users are visiting are legitimate, the monitor uses the DOM interface to retrieve the data, as in those cases client scripts manipulation is not an issue. However, in detection mode, when the user is interacting with websites that may be phishing websites, client scripts manipulation is a real threat. The monitor takes hence a

different approach. The monitor performs like a keylogger, it collects user data by listening to a user's keystrokes for each element in the webpage. Such key logging is very efficient to run and since UBPD runs in detection mode only occasionally, there is little performance overhead.

B. Activation of the detection engine & denial of service attack

The detection engine must be activated before the data the user entered has been sent to a remote server. Normally webpages forward the user submitted data using built-in functions, such as the one provided by the standard web form (Almost every legitimate website uses this function to submit the data). So the detection engine is triggered and data submission is suspended when the monitor discovers the use of such functions. This is achieved by listening for 'DOMActivate' [16] events. These events are fired by the browser when such built-in functions have been used. This is the mechanism the monitor uses to activate the detection engine when the system is running in training mode.

However, phishers can use client side scripts to implement these built in functions. For example, by using Javascript they can access the DOM interface to retrieve the data users entered and use AJAX (Asynchronous Javascript And XML) techniques to send the data to the server before the user clicks the submit button. To prevent this evasion, UBPD could monitor the client script function calls when it is running under the detection mode. The function calls should be monitored are 'open()' and 'send()' from the 'xmlhttprequest' API [36]. These are the only two functions that client scripts can use to send data to the server. Once the monitor discovers such function calls, the function is suspended and the detection engine is triggered. Regardless of the data the client scripts would like to send, the detection engine always works on all the data the user has entered to the current webpage. The function call is only resumed if the detection engine thinks the current website is legitimate.

Since the detection engine would be triggered once the 'xmlhttprequest' API has been called, the attackers can issue this function call continuously to freeze the browser. To prevent this, the monitor can keep a record of whether there is user input since the last time the detection engine was activated for the same webpage. If there is then the detection engine will be activated, otherwise, not.

We plan to implement such client side scripts function call detection by using the techniques in Venkman (a Firefox JavaScript Debugger Addon) in the next version of UBPD. Nevertheless, if attacked by this evasion technique, at least UBPD would still be able to inform users what they have done, and preserve the evidence for possible analysis. One of the major problem for users is they do not know when they have been attacked and how. It gives users a good chance to reduce the damage caused by the phishing attacks.

V. EVALUATION

We have carried out two experiments to evaluate the effectiveness of UBPD in terms of the two following rates:

TABLE I
CHARACTERISTICS OF USER PROFILE

	Alice	Bob	Carol	Dave
Reuse	No	No	Yes	Yes
Uniqueness	Strong	Weak	Strong	Weak

- **False negative:** The system fails to recognise a phishing attack.
- **False positive:** The system recognises a legitimate website as a phishing website.

In addition we also search for a useful default threshold value (mentioned in section III-B4). For both experiments UBPD was modified to not present the warning dialogue, instead it records the phishing score results as well as the URLs for later analysis.

A. False Negative Rate

From PhishTank[32] and millersmiles[25] we randomly collected 463 recently reported phishing webpages, which target Ebay, Paypal, and Natwest bank. We created four user profiles, which describe four artificial users' binding relationships with the three targeted websites. The four user profiles have different characteristics as shown in Table I. 'Reuse' indicates maximum possible reuse of authentication credentials. In this case the user would have same user name and password for Ebay and Paypal. 'Uniqueness' indicates whether the user would use the exact data they shared with a legitimate website at other places. For example if Bob chooses his email address as password then the uniqueness is weak, because Bob is very likely to tell other websites his email address. If Bob uses some random string as his password, then the uniqueness is strong, because this random string is unlikely to be used with any other websites.

We entered the artificial authentication credentials to each phishing webpages. Regardless of the characteristics of the user profile, the detection result is the same for all four users: 459 pages had a phishing score of 1, and 4 had a phishing score of 0. Thus only four pages evaded detection – a 99.14 percent detection rate. Compared to other existing phishing website detection systems, UBPD's detection rate may not be significantly better. **Its biggest advantage** is that its detection method detects essential characteristics of a phishing attack, namely that phishing web pages request authentication credentials. The details of how users may be manipulated may change with future phishing attacks, but the requesting of such details remains constant. Other detection systems based on the analysis of incoming data will need to adapt and be redesigned for future phishing attacks; UBPD will not.

Detailed analysis confirms that the detection result is determined mainly by the information requested by the phishing webpage. Table II shows the classification of the phishing webpages based on the type of information they requested. 92% of the collected phishing webpages asked only for authentication credentials and 7.14% of the collected phishing webpages asked both for personal and authentication credentials.

TABLE II
STATISTICS OF PHISHING WEBSITES BASED ON THE TYPE OF INFORMATION REQUESTED. AC: AUTHENTICATION CREDENTIALS; PI: PERSONAL INFORMATION

	Ebay	Paypal	Natwest
AC	211	176	39
AC+PI	22	5	6
PI	4	0	0
Total	237	181	45

The four phishing web pages that UBPD failed to detect asked only for personal information such as full name, address, telephone number and mother's maiden name. In fact, they can not be detected by UBPD no matter what the threshold value is. However, it is impractical for phishing attacks to ask for personal information without requesting authentication credentials first, because those phishing webpages are not presented to users when a user victim first arrives at the phishing website. Those phishing websites would normally first present the user with a login webpage before directing the user to the webpage that asking for the personal information (none of the four phishing webpages were the landing page of the phishing attacks). Otherwise such practice would be seemed abnormal, and make potential victims very suspicious. As a result, UBPD can detect the phishing attacks and stop users from even reaching the phishing webpages that ask for personal information.

The sample size in this experiment is large and we might have some expectation that this would be reasonably indicative of success rate when deployed 'in the wild'.

B. False Positive Rate

Five volunteers were provided with the information needed to install UBPD on their machine. We did not explicitly ask them to train UBPD with all their binding relationships, because we wanted to see how users would train UBPD and what the false positives would be in reality if the user has not properly trained UBPD. At the end of one week, we collected the result log from their machines.

The volunteers were three male and two female science students. They all used Firefox as their main web browser. They were all regular Internet users (in average over three hours per day). As a result the UBPD was activated a large number of times and the interactions that occurred during the experiments covered a wide range of types of interaction. Another reason we chose those volunteers is because they are the most unlikely user group to fall victims to phishing attacks [17] and so we can safely assume they have all interacted with legitimate websites. In total the volunteers interacted with 76 distinct websites, submitted data to those websites 2107 times, and UBPD ran in detection mode only 81 times. In fact all the websites volunteers visited were legitimate. On 59 occasions the phishing score was 0, on five interactions gave a score of 0.25, on 13 occasions the score was 0.5, and the score was 1 on three occasions.

The phishing score was 1 when users interacted with three legitimate websites (the registration webpages of videojug.com

and surveys.com, and the authentication webpage of a web forum). We asked the volunteers what data they supplied to those webpages. It seems that the reuse of authentication credentials on creating new accounts is the reason. In this experiment, the warning dialog is not presented, as we did not aim to test usability. But if it does, then the user must make decision to train UBPD to remember these new binding relationships. To avoid the user's confusion about what is the right choice when the warning dialog is presented, the dialog always reminds the user of the legitimate websites UBPD is aware of, and tells the user that if the user is sure the current website is legitimate, and the website is not remembered by UBPD, then they need to update their binding relationships (see the figure 4 in section III-B6). This requires no technical knowledge and should be quite easy to understand. There are only two choices provided by the dialog: to update the profile and submit the data; or do not send the user submitted data and close the phishing webpage. There is no third choice provided by the dialog, in this way we force the user to make the security decision and they can not just ignore the warnings given by the system.

Many websites force users to supply an email address as the user name. As a result, the user's email address is kept in the user profile as part of user's authentication credentials. This email address almost inevitably will be given out to other websites, which are not contained in the user profile, for various reasons such as contact method, activate the newly registered account, etc. Thus even when the user does not intend to give out their credentials, the email address nevertheless is shared and our system simply confirmed that by calculating the phishing score of 0.5 (which means half of the data the user has shared with a legitimate website was given away to a website that was not in user's profile) on 13 occasions.

For one volunteer five interactions gave a phishing score of 0.25. The user had an account at a major bank, the authentication credentials for which compromised four data items. One of these was the family name. For other sites not included in the user's profile asking for this information caused our system to identify the sharing of the data.

Based on the figures from both experiments we decided to set the default threshold value to 0.6. First, it can successfully detect phishing webpages asking for more than half of the credentials the user has shared with a legitimate website (99.14% of the phishing websites in Experiment One can be detected). It also generated few false positives. The false positive rate of the system is 0.0014 (obtained by dividing the number of false positives generated with the total number of times the UBPD was activated).

Of course this is only a preliminary false positive evaluation, but it still shows the system has a small false positive rate, and it also shows that the reuse of the authentication credentials and partial training are the main cause of the false positives.

VI. DISCUSSION

A. Why UBPD is useful

Besides its high detection accuracy, UBPD is useful also because it complements existing detection systems. First UBPD detects phishing websites based on users' behaviours, not the incoming data that attackers can manipulate freely. Violation of the binding relationships cannot be changed no matter what techniques phishers choose to use. As the evaluation proves, UBPD is able to consistently detect phishing webpages regardless of how they are implemented as long as they ask for authentication credentials. In contrast detection systems based on incoming data may find it difficult to deal with novel and sophisticated spoofing techniques. UBPD analyses the identity of websites using both IP addresses and domain names, it can detect pharming attacks, which are undetectable by many existing systems. Being independent of the incoming data means low cost in maintenance, the system does not need updating when attackers vary their techniques, and so we have far fewer evasion techniques to deal with.

Some systems that try to stop phishing attacks from reaching the users (phishing site take down, botnet take down, phishing email filter, etc.), some have tried to detect phishing webpages as soon as users arrive at a new web page (Phishing URL blacklist, netcraft toolbar, spoofguard, CARTINA, etc), and some have tried to provide useful information to help users to detect phishing websites. However, there are no other systems that work at the stage when phishing webpages have somehow penetrated through and users have started to give out information to them. Thus our system plugs a major remaining gap, and it can be easily combined with other phishing detection systems to provide multi-stage protection.

B. Performance

The two working modes of UBPD reduce the computing complexity. Computing in detection mode is more expensive, but it runs only occasionally (in our second experiment UBPD only in detection mode 81 out of 2107 times). Computing in the detection mode is still light weight for the computing power of an average personal computer (none of the volunteers noticed any delay). As a result, UBPD is efficient to run and adds few delay to the existing user-webpage interaction experience.

C. Limitations and Future Work

UBPD should be viewed as an initial but promising effort towards detecting phishing by analysing user behaviour. Despite its detection effectiveness, there are some limitations within the implementation and design.

Currently UBPD cannot handle all types of authentication credentials. It can handle static type authentication credentials such as user name, password, security questions, etc, but dynamic authentication credentials shared between users and the legitimate websites cannot be handled by the current implementation (e.g. one-time passwords). In future we will investigate how to handle such credentials.

There is also another method for a phisher to evade detection. Attackers might compromise a website within a user's personal whitelist and host the phishing website under the website's domain (perhaps for several hours to a couple of days). This method is difficult to carry out, and it is unlikely to happen if there are still easier ways to launch phishing attacks.

Detection itself requires little computing time, in fact, retrieving the data from WHOIS database is the most time consuming part. It depends on the type of network the user is using and the workload on the WHOIS database. In future we could add some cache functionality to reduce the number of queries for each detection. Currently the system produces a slight delay because of the WHOIS database query.

Although we have paid attention to the usability design of the system, we have not systematically evaluated it. Further user evaluation will inform development.

VII. ACKNOWLEDGEMENT

I wish to thank all the people helped us, and especially the Engineering and Physical Sciences Research Council (EPSRC) of United Kingdom, without their sponsorship of the project "Defending the weakest link: Intrusion via Social Engineering" the detection method would not be invented.

VIII. CONCLUSION

We have used a novel paradigm—analysis of the users' behaviours—to detect phishing websites. We have shown that it is an accurate method, discussed how it has been designed and implemented to be hard to circumvent, and have discussed its unique strength in protecting users from phishing threats. UBPD is not designed to replace existing techniques. Rather it should be used to complement other techniques, to provide better overall protection. We believe our approach fills a significant gap in current anti-phishing technology capability.

REFERENCES

- [1] <http://home.versatel.nl/MAvanEverdingen/Code/>.
- [2] Spamassassin's official website. <http://spamassassin.apache.org/>.
- [3] Anti-phishing work group home page, 2007. <http://www.antiphishing.org/>.
- [4] C. Abad. The economy of phishing: A survey of the operations of the phishing market. *First Monday*, 10(9), 2005.
- [5] K. Albrecht, N. Burri, and R. Wattenhofer. Spamoto—An Extendable Spam Filter System. In *2nd Conference on Email and Anti-Spam (CEAS)*, Stanford University, Palo Alto, California, USA, July 2005.
- [6] P. Behera and N. Agarwal. A confidence model for web browsing. In *Toward a More Secure Web—W3C Workshop on Transparency and Usability of Web Authentication*, 2006.
- [7] M. Chandrasekaran, R. Chinchain, and S. Upadhyaya. Mimicking user response to prevent phishing attacks. In *IEEE International Symposium on a World of Wireless, Mobile, and Multimedia networks*, 2006.
- [8] N. Chou, R. Ledesma, Y. Teraguchi, D. Boneh, and J. C. Mitchell. Client-side defense against web-based identity theft. In *NDSS '04: Proceedings of the 11th Annual Network and Distributed System Security Symposium*, February 2004.
- [9] R. Dhamija, D. Tygar, and M. Hearst. Why phishing works. *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, ACM Special Interest Group on Computer-Human Interaction:581–590, 2006.
- [10] X. Dong, J. A. Clark, and J. Jacob. A user-phishing interaction model. In *Conference on Human System Interaction*, 2008.
- [11] I. Fette, N. Sadeh, and A. Tomasic. Learning to detect phishing emails. In *WWW '07: Proceedings of the 16th international conference on World Wide Web*, pages 649–656, New York, NY, USA, 2007. ACM Press.
- [12] D. Florencio and C. Herley. A large-scale study of web password habits. In *WWW '07: Proceedings of the 16th international conference on World Wide Web*, pages 657–666, New York, NY, USA, 2007. ACM Press.
- [13] R. Franco. Better website identification and extended validation certificates in ie7 and other browsers. IEBlog, November 2005.
- [14] B. Goodger, I. Hickson, D. Hyatt, and C. Waterson. Xml user interface language (xul) 1.0. Technical report, Mozilla Org., 2001.
- [15] S. Hartman. Ietf-draft: Requirements for web authentication resistant to phishing. Technical report, MIT, 2007.
- [16] B. Hhrmann, P. L. Hgaret, and T. Pixley. Document object model events. Technical report, W3C, 2007.
- [17] T. Jagatic, N. Johnson, M. Jakobsson, and F. Menczer. Social phishing. *ACM Communication*, October 2007.
- [18] M. Jakobsson. Modeling and preventing phishing attacks. In *Phishing Panel in Financial Cryptography '05*, 2005.
- [19] M. Jakobsson. Human factors in phishing. *Privacy & Security of Consumer Information '07*, 2007.
- [20] M. Jakobsson, A. Tsow, A. Shah, E. Blevis, and Y.-K. Lim. What instills trust? a qualitative study of phishing. In *Extended abstract, USEC '07*, 2007.
- [21] P. A. Johnston. <http://pajhome.org.uk/crypt/index.html>.
- [22] A. Litan. Toolkit: E-commerce loses big because of security concerns. Technical report, Garnter Research, 2006.
- [23] T. McCall. Gartner survey shows phishing attacks escalated in 2007; more than \$3 billion lost to these attacks. Technical report, Gartner Research, 2007.
- [24] Microsoft. Anti-phishing white paper. Technical report, Microsoft, 2005.
- [25] MillerSmiles. Official website. <http://www.millersmiles.co.uk>.
- [26] T. Moore and R. Clayton. Examining the impact of website take-down on phishing. In *eCrime '07: Proceedings of the anti-phishing working groups 2nd annual eCrime researchers summit*, pages 1–13, New York, NY, USA, 2007. ACM.
- [27] Mozilla. Phishing protection, 2007. <http://www.mozilla.com/en-US/firefox/phishing-protection/>.
- [28] Netcraft, 2007. <http://toolbar.netcraft.com/>.
- [29] G. Ollmann. The phishing guide. Technical report, NGSS.
- [30] G. Ollmann. The pharming guide. Technical report, Next Generation Security Software Ltd., 2005.
- [31] Y. Pan and X. Ding. Anomaly based web phishing page detection. *acsac*, 0:381–392, 2006.
- [32] Phishtank, 2007. <http://www.phishtank.com/>.
- [33] S. Schechter, R. Dhamija, A. Ozment, and I. Fischer. The emperor's new security indicators: An evaluation of website authentication and the effect of role playing on usability studies. In *2007 IEEE Symposium on Security and Privacy*, 2007.
- [34] R. Security. Enhancing one-time passwords for protection against real-time phishing attacks. Technical report, RSA, 2007.
- [35] G. Staikos. Web browser developers work together on security. Web, November 2005.
- [36] A. van Kesteren. The xmlhttprequest objec. Technical report, W3C, 2006.
- [37] W3C. Web security context—working group charter. Web, 2006.
- [38] D. Watson, T. Holz, and S. Mueller. Know your enemy: Phishing. Technical report, The Honeynet Project & Research Alliance, 2005.
- [39] M. Wu, R. C. Miller, and S. L. Garfinkel. Do security toolbars actually prevent phishing attacks? In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 601–610, New York, NY, USA, 2006. ACM Press.
- [40] M. Wu, R. C. Miller, and G. Little. Web wallet: preventing phishing attacks by revealing user intentions. pages 102–113, 2006.
- [41] Y. Zhang, J. I. Hong, and L. F. Cranor. Cantina: a content-based approach to detecting phishing web sites. In *WWW '07: Proceedings of the 16th international conference on World Wide Web*, pages 639–648, New York, NY, USA, 2007. ACM Press.