

Online Detection and Prevention of Phishing Attacks (Invited Paper)

Juan Chen

Institute of Communications Engineering
Nanjing 210007, P.R. China
icechj@msn.com

Chuanxiong Guo

Institute of Communications Engineering
Nanjing 210007, P.R. China
xguo@ieee.org

Abstract—Phishing is a new type of network attack where the attacker creates a replica of an existing Web page to fool users (e.g., by using specially designed e-mails or instant messages) into submitting personal, financial, or password data to what they think is their service provider's Web site. In this paper, we propose a new end-host based anti-phishing algorithm, which we call LinkGuard, by utilizing the generic characteristics of the hyperlinks in phishing attacks. These characteristics are derived by analyzing the phishing data archive provided by the Anti-Phishing Working Group (APWG). Because it is based on the generic characteristics of phishing attacks, LinkGuard can detect not only known but also unknown phishing attacks. We have implemented LinkGuard in Windows XP. Our experiments verified that LinkGuard is effective to detect and prevent both known and unknown phishing attacks with minimal false negatives. LinkGuard successfully detects 195 out of the 203 phishing attacks. Our experiments also showed that LinkGuard is light-weighted and can detect and prevent phishing attacks in real-time.

Index Terms—Network security, Phishing attacks, Hyperlink, LinkGuard algorithm.

I. INTRODUCTION

The word 'Phishing' initially emerged in 1990s. The early hackers often use 'ph' to replace 'f' to produce new words in the hacker's community, since they usually hack by phones. Phishing is a new word produced from 'fishing', it refers to the act that the attacker allure users to visit a faked Web site by sending them faked e-mails (or instant messages), and stealthily get victim's personal information such as user name, password, and national security ID, etc. These information then can be used for future target advertisements or even identity theft attacks (e.g., transfer money from victims' bank account).

The frequently used attack method is to send e-mails to potential victims, which seemed to be sent by banks, online organizations, or ISPs. In these e-mails, they will make up some causes, e.g. the password of your credit card had been mis-entered for many times, or they are providing upgrading services, to allure you visit their Web site to conform or modify your account number and password through the hyperlink provided in the e-mail. You will then be linked to a counterfeited Web site after clicking those links. The style, the functions performed, sometimes even the URL of these faked Web sites

are similar to the real Web site. It's very difficult for you to know that you are actually visiting a malicious site. If you input the account number and password, the attackers then successfully collect the information at the server side, and is able to perform their next step actions with that information (e.g., withdraw money out from your account).

Phishing itself is not a new concept, but it's increasingly used by phishers to steal user information and perform business crime in recent years. Within one to two years, the number of phishing attacks increased dramatically. According to Gartner Inc., for the 12 months ending April 2004, "there were 1.8 million phishing attack victims, and the fraud incurred by phishing victims totaled \$1.2 billion" [6].

According to the statistics provided by the Anti-Phishing Working Group (APWG) [2], in March 2006, the total number of unique phishing reports submitted to the APWG was 18,480; and the top three phishing site hosting countries are, the United States (35.13%), China (11.93%), and the Republic of Korea (8.85%). The infamous phishing attacks happened in China in recent years include the events to counterfeit the Bank of China (real Web site www.bank-of-china.com, counterfeited Web site www.bank-off-china.com), the Industrial and Commercial Bank of China (real Web site www.icbc.com.cn, faked web site www.lcbc.com.cn), the Agricultural Bank of China (real webs ite www.95599.com, faked Web site www.965555.com), etc.

In this paper, we study the common procedure of phishing attacks and review possible anti-phishing approaches. We then focus on end-host based anti-phishing approach. We first analyze the common characteristics of the hyperlinks in phishing e-mails. Our analysis identifies that the phishing hyperlinks share one or more characteristics as listed below: 1) the visual link and the actual link are not the same; 2) the attackers often use dotted decimal IP address instead of DNS name; 3) special tricks are used to encode the hyperlinks maliciously; 4) the attackers often use fake DNS names that are similar (but not identical) with the target Web site. We then propose an end-host based anti-phishing algorithm which we call LinkGuard, based on the characteristics of the phishing hyperlink. Since LinkGuard is character-based, it can detect and prevent not only known phishing attacks but also unknown ones. We have implemented LinkGuard in Windows XP, and our experiments indicate that LinkGuard is light-weighted in

that it consumes very little memory and CPU circles, and most importantly, it is very effective in detecting phishing attacks with minimal false negatives. LinkGuard detects 195 attacks out of the 203 phishing archives provided by APWG without knowing any signatures of the attacks.

The rest of this paper is organized as follows. In Section II, we give the general procedure of a phishing attack and provide the available methods to prevent phishing attacks. We then analyze the characteristics of the hyperlinks used in phishing attacks and present the LinkGuard algorithm in Section III. Section IV describes our implementation of the LinkGuard system and gives the experimental results. Section V concludes this paper.

II. PHISHING ATTACK PROCEDURE AND PREVENTION METHODS

In this paper, we assume that phishers use e-mail as their major method to carry out phishing attacks. Nonetheless, our analysis and algorithm can be applied to attacks that use other means such as instant messaging.

A. The Procedure of Phishing Attacks

In general, phishing attacks are performed with the following four steps:

- 1) Phishers set up a counterfeited Web site which looks exactly like the legitimate Web site, including setting up the web server, applying the DNS server name, and creating the web pages similar to the destination Web site, etc.
- 2) Send large amount of spoofed e-mails to target users in the name of those legitimate companies and organizations, trying to convince the potential victims to visit their Web sites.
- 3) Receivers receive the e-mail, open it, click the spoofed hyperlink in the e-mail, and input the required information.
- 4) Phishers steal the personal information and perform their fraud such as transferring money from the victims' account.

B. Approaches to Prevent Phishing Attacks

There are several (technical or non-technical) ways to prevent phishing attacks: 1) educate users to understand how phishing attacks work and be alert when phishing-alike e-mails are received; 2) use legal methods to punish phishing attackers; 3) use technical methods to stop phishing attackers. In this paper, we only focus on the third one.

Technically, if we can cut off one or several of the steps that needed by a phishing attack, we then successfully prevent that attack. In what follows, we briefly review these approaches.

1) *Detect and block the phishing Web sites in time:* If we can detect the phishing Web sites in time, we then can block the sites and prevent phishing attacks. It's relatively easy to (manually) determine whether a site is a phishing site or not, but it's difficult to find those phishing sites out in time. Here we list two methods for phishing site detection. 1) The Web

master of a legal Web site periodically scans the root DNS for suspicious sites (e.g. www.1cbbc.com.cn vs. www.icbc.com.cn).

2) Since the phisher must duplicate the content of the target site, he must use tools to (automatically) download the Web pages from the target site. It is therefore possible to detect this kind of download at the Web server and trace back to the phisher. Both approaches have shortcomings. For DNS scanning, it increases the overhead of the DNS systems and may cause problem for normal DNS queries, and furthermore, many phishing attacks simply do not require a DNS name. For phishing download detection, clever phishers may easily write tools which can mimic the behavior of human beings to defeat the detection.

2) *Enhance the security of the web sites:* The business Web sites such as the Web sites of banks can take new methods to guarantee the security of users' personal information. One method to enhance the security is to use hardware devices. For example, the Barclays bank provides a hand-held card reader to the users. Before shopping in the net, users need to insert their credit card into the card reader, and input their (personal identification number) PIN code, then the card reader will produce a onetime security password, users can perform transactions only after the right password is input [12]. Another method is to use the biometrics characteristic (e.g. voice, fingerprint, iris, etc.) for user authentication. For example, Paypal had tried to replace the single password verification by voice recognition to enhance the security of the Web site. With these methods, the phishers cannot accomplish their tasks even after they have gotten part of the victims' information. However, all these techniques need additional hardware to realize the authentication between the users and the Web sites, hence will increase the cost and bring certain inconvenience. Therefore, it still needs time for these techniques to be widely adopted.

3) *Block the phishing e-mails by various spam filters:* Phishers generally use e-mails as 'bait' to allure potential victims. SMTP (Simple Mail Transfer Protocol) [11] is the protocol to deliver e-mails in the Internet. It is a very simple protocol which lacks necessary authentication mechanisms. Information related to sender, such as the name and email address of the sender, route of the message, etc., can be counterfeited in SMTP. Thus, the attackers can send out large amounts of spoofed e-mails which are seemed from legitimate organizations.

The phishers hide their identities when sending the spoofed e-mails, therefore, if anti-spam systems can determine whether an e-mail is sent by the announced sender (Am I Whom I Say I Am?), the phishing attacks will be decreased dramatically. From this point, the techniques that preventing senders from counterfeiting their Send ID (e.g. SIDF of Microsoft [8]) can defeat phishing attacks efficiently.

SIDF is a combination of Microsoft's Caller ID for E-mail and the SPF (Sender Policy Framework) [13] developed by Meng Weng Wong. Both Caller ID and SPF check e-mail sender's domain name to verify if the e-mail is sent from a server that is authorized to send e-mails of that domain, and

from that to determine whether that e-mail use spoofed e-mail address. If it's faked, the Internet service provider can then determine that e-mail is a spam e-mail.

The spoofed e-mails used by phishers are one type of spam e-mails. From this point of view, the spam filters [1], [4] can also be used to filter those phishing e-mails. For example, blacklist, whitelist, keyword filters, Bayesian filters with self-learning abilities, and E-Mail Stamp, etc., can all be used at the e-mail server or client systems. Most of these anti-spam techniques perform filtering at the receiving side by scanning the contents and the address of the received e-mails. And they all have pros and cons as discussed below. Blacklist and whitelist cannot work if the names of the spammers are not known in advance. Keyword filter and Bayesian filters can detect spam based on content, hence can detect unknown spasm. But they can also result in false positives and false negatives. Furthermore, spam filters are designed for general spam e-mails and may not very suitable for filtering phishing e-mails since they generally do not consider the specific characteristics of phishing attacks.

4) *Install online anti-phishing software in user's computers:* Despite all the above efforts, it is still possible for the users to visit the spoofed Web sites. As a last defense, users can install anti-phishing tools in their computers. The anti-phishing tools in use today can be divided into two categories: blacklist/whitelist based and rule-based.

- Category I: When a user visits a Web site, the anti-phishing tool searches the address of that site in a blacklist stored in the database. If the visited site is on the list, the anti-phishing tool then warns the users. Tools in this category include ScamBlocker from the EarthLink company [5], PhishGuard [10], and Netcraft [9], etc. Though the developers of these tools all announced that they can update the blacklist in time, they cannot prevent the attacks from the newly emerged (unknown) phishing sites.
- Category II: this category of tools uses certain rules in their software, and checks the security of a Web site according to these rules. Examples of this type of tools include SpoofGuard developed by Stanford [3], TrustWatch of the GeoTrust [7], etc. SpoofGuard checks the domain name, URL (includes the port number) of a Web site, it also checks whether the browser is directed to the current URL via the links in the contents of e-mails. If it finds that the domain name of the visited Web site is similar to a well-known domain name, or if they are not using the standard port, SpoofGuard will warn the users. In TrustWatch, the security of a Web site is determined by whether it has been reviewed by an independent trusted third party organization. Both SpoofGuard and TrustWatch provide a toolbar in the browsers to notify their users whether the Web site is verified and trusted.

It is easy to observe that all the above defense methods are useful and complementary to each other, but none of them are perfect at the current stage. In the rest of the paper, we focus on end-host based approach and propose an end-

host based LinkGuard algorithm for phishing detection and prevention. To this end, our work follows the same approach as [3]. Our work differs from [3] in that: 1) LinkGuard is based on our careful analysis of the characteristics of phishing hyperlinks whereas SpoofGuard is more like a framework; 2) LinkGuard has a verified very low false negative rate for unknown phishing attacks whereas the false negative property of SpoofGuard is still not known. In next section, we first study the characteristics of the hyperlinks in phishing e-mails and then we propose the LinkGuard algorithm.

III. LINKGUARD

A. Classification of the hyperlinks in the phishing e-mails

In order to (illegally) collect useful information from potential victims, phishers generally tries to convince the users to click the hyperlink embedded in the phishing e-mail. A hyperlink has a structure as follows.

```
<a href="URI"> Anchor text <\a>
```

where 'URI' (universal resource identifiers) provides the necessary information needed for the user to access the networked resource and 'Anchor text' is the text that will be displayed in user's Web browser. Examples of URIs are <http://www.google.com>, <https://www.icbc.com.cn/login.html>, <ftp://61.112.1.90:2345>, etc. 'Anchor text' in general is used to display information related to the URI to help the user to better understand the resources provided by the hyperlink. In the following hyperlink, the URI links to the phishing archives provided by the APWG group, and its anchor text "Phishing Archive" informs the user what's the hyperlink is about.

```
<a href="http://www.antiphishing.org/phishing_archive.html"> Phishing Archive </a>
```

Note that the content of the URI will not be displayed in user's Web browser. Phishers therefore can utilize this fact to play trick in their 'bait' e-mails. In the rest of the paper, we call the URI in the hyperlink the actual link and the anchor text the visual link.

After analyzing the 203 (there are altogether 210 phishing e-mails, with 7 of them with incomplete information or with malware attachment and do not have hyperlinks) phishing e-mail archives from Sep. 21st 2003 to July 4th 2005 provided by APWG [6]. We classified the hyperlinks used in the phishing e-mail into the following categories:

- 1) The hyperlink provides DNS domain names in the anchor text, but the destination DNS name in the visible link doesn't match that in the actual link. For instance, the following hyperlink:

```
<a href = "http://www.profusenet.net/checksession.php"> https://secure.regionset.com/EBanking/logon/</a>
```

appears to be linked to secure.regionset.com, which is the portal of a bank, but it actually is linked to a phishing site www.profusenet.net.
- 2) Dotted decimal IP address is used directly in the URI or the anchor text instead of DNS name. See below for an example.

 SIGN IN

- 3) The hyperlink is counterfeited maliciously by using certain encoding schemes. There are two cases: a) The link is formed by encoding alphabets into their corresponding ASCII codes. See below for such a hyperlink.

 www.citibank.com

while this link is seemed pointed www.citibank.com, it actually points to http://4.34.195.41:341/index.htm.

b) Special characters (e.g. @ in the visible link) are used to fool the user to believe that the e-mail is from a trusted sender. For instance, the following link seems is linked to amazon, but it actually is linked to IP address 69.10.142.34.

http://www.amazon.com:fvthsgbljhfc83infoupdate @69.10.142.34.

- 4) The hyperlink does not provide destination information in its anchor text and uses DNS names in its URI. The DNS name in the URI usually is similar with a famous company or organization. For instance, the following link seems to be sent from paypal, but it actually is not. Since paypal-cgi is actually registered by the phisher to let the users believe that it has something to do with paypal

 Click here to confirm your account

- 5) The attackers utilize the vulnerabilities of the target Web site to redirect users to their phishing sites or to launch CSS (cross site scripting) attacks. For example, the following link

 Click here <a>

Once clicked, will redirect the user to the phishing site 200.251.251.10 due to a vulnerability of usa.visa.com.

Table 1 summarizes the number of hyperlinks and their percentages for all the categories. It can be observed that most of the phishing e-mails use faked DNS names (category 1, 44.33%) or dotted decimal IP addresses (category 2, 41.87%). Encoding tricks are also frequently used (category 3a and 3b, 17.24%). And phishing attackers often try to fool users by setting up DNS names that are very similar with the real e-commerce sites or by not providing destination information in the anchor text (category 4). Phishing attacks that utilize the vulnerability of Web sites (category 5) are of small number (2%) and we leave this type of attacks for future study.

Note that a phishing hyperlink can belong to several categories at the same time. For instance, an attacker may use tricks from both categories 1 and 3 at the same time to increase his success chance. Hence the sum of percentages is larger than 1.

Category	Number of links	Percentage
1	90	44.33%
2	85	41.87%
3.a	19	9.36%
3.b	16	7.88%
4	67	33%
5	4	2%

TABLE I

THE CATEGORIES OF HYPERLINKS IN PHISHING E-MAILS.

Once the characteristics of the phishing hyperlinks are understood, we are able to design anti-phishing algorithms that can detect known or unknown phishing attacks in real-time. We present our LinkGuard algorithm in the next subsection.

B. The LinkGuard algorithm

LinkGuard works by analyzing the differences between the visual link and the actual link. It also calculates the similarities of a URI with a known trusted site. The algorithm is illustrated in Fig. 1. The following terminologies are used in the algorithm.

v_link: visual link;
a_link: actual_link;
v_dns: visual DNS name;
a_dns: actual DNS name;
sender_dns: sender’s DNS name.

```
int LinkGuard(v_link, a_link) {
1  v_dns = GetDNSName(v_link);
2  a_dns = GetDNSName(a_link);
3  if ((v_dns and a_dns are not
4    empty) and (v_dns != a_dns))
5    return PHISHING;
6  if (a_dns is dotted decimal)
7    return POSSIBLE_PHISHING;
8  if(a_link or v_link is encoded)
9  {
10   v_link2 = decode (v_link);
11   a_link2 = decode (a_link);
12   return LinkGuard(v_link2, a_link2);
13 }
14 /* analyze the domain name for
15   possible phishing */
16 if(v_dns is NULL)
17   return AnalyzedDNS(a_link);
}
```

Fig. 1. Description of the LinkGuard algorithm.

The LinkGuard algorithm works as follows. In its main routine *LinkGuard*, it first extracts the DNS names from the actual and the visual links (lines 1 and 2). It then compares the actual and visual DNS names, if these names are not the same, then it is phishing of category 1 (lines 3-5). If dotted decimal IP address is directly used in actual_dns , it is then a possible phishing attack of category 2 (lines 6 and 7). We will delay the discussion of how to handle possible phishing attacks later. If the actual link or the visual link is encoded

```

int AnalyzeDNS (actual_link) {
    /* Analyze the actual DNS name according
       to the blacklist and whitelist*/
    18 if (actual_dns in blacklist)
    19     return PHISHING;
    20 if (actual_dns in whitelist)
    21     return NOTPHISHING;
    22 return PatternMatching(actual_link);
}
int PatternMatching(actual_link){
    23 if (sender_dns and actual_dns are different)
    24     return POSSIBLE_PHISHING;
    25 for (each item prev_dns in seed_set)
    26 {
    27     bv = Similarity(prev_dns, actual_link);
    28     if (bv == true)
    29         return POSSIBLE_PHISHING;
    30 }
    31 return NO_PHISHING;
}
float Similarity (str, actual_link) {
    32 if (str is part of actual_link)
    33     return true;
    34 int maxlen = the maximum string
    35     lengths of str and actual_dns;
    36 int minchange = the minimum number of
    37     changes needed to transform str
    38     to actual_dns (or vice verse);
    39 if (thresh<(maxlen-minchange)/maxlen<1)
    40     return true
    41 return false;
}

```

Fig. 2. The subroutines used in the LinkGuard algorithm.

(categories 3 and 4), we first decode the links, then recursively call *LinkGuard* to return a result (lines 8-13). When there is no destination information (DNS name or dotted IP address) in the visual link (category 5), LinkGuard calls *AnalyzeDNS* to analyze the actual_dns (lines 16 and 17). LinkGuard therefore handles all the 5 categories of phishing attacks.

AnalyzeDNS and the related subroutines are depicted in Fig. 2. In *AnalyzeDNS*, if the actual_dns name is contained in the blacklist, then we are sure that it is a phishing attack (lines 18 and 19). Similarly, if the actual_dns is contained in the whitelist, it is therefore not a phishing attack (lines 20 and 21). If the actual_dns is not contained in either whitelist or blacklist, *PatternMatching* is then invoked (line 22).

PatternMatching is designed to handle unknown attacks (blacklist/whitelist is useless in this case). For category 5 of the phishing attacks, all the information we have is the actual_link from the hyperlink (since the visual link does not contain DNS or IP address of the destination site), which provide very little information for further analysis. In order to resolve this problem, we try two methods: First, we extract the sender e-mail address from the e-mail. Since phishers generally try to fool users by using (spoofed) legal DNS names in the sender e-mail address, we expect that the DNS name in the sender address will be different from that in the actual link. Second, we proactively collect DNS names that are manually input by

the user when she surfs the Internet and store the names into a *seed_set*, and since these names are input by the user by hand, we assume that these names are trustworthy. *PatternMatching* then checks if the actual DNS name of a hyperlink is different from the DNS name in the sender's address (lines 23 and 24), and if it is quite similar (but not identical) with one or more names in the *seed_set* by invoking the *Similarity* (lines 25-30) procedure.

Similarity checks the maximum likelihood of actual_dns and the DNS names in *seed_set*. As depicted in Fig. 2, the similarity index between two strings are determined by calculating the minimal number of changes (including insertion, deletion, or revision of a character in the string) needed to transform a string to the other string. If the number of changes is 0, then the two strings are identical; if the number of changes is small, then they are of high similarity; otherwise, they are of low similarity. For example, the similarity index of 'microsoft' and 'micr0s0ft' is 7/9 (since we need change the 2 '0's in micr0s0ft to 'o'). Similarly, the similarity index of 'paypal' and 'paypal-cgi' is 6/10 (since we need to remove the last 4 chars from paypal-cgi), and the similarity index of '95559' and '955559' is 5/6 (since we need to insert a '5' to change '95559' to '955559').

If the two DNS names are similar but not identical, then it is a possible phishing attack. For instance, *PatternMatching* can easily detect the difference between www.icbc.com.cn (which is a good e-commerce Web site) and www.lcbc.com.cn (which is a phishing site), which has similarity index 75%. Note that *PatternMatching* may treat www.lcbc.com.cn as a normal site if the user had never visit www.lcbc.com.cn before. This false negative, however, is unlikely to cause any severe privacy or financial lose to the user, since she actually does not have anything to lose regarding the Web site www.icbc.com.cn (since she never visits that Web site before)!

C. False positives and false negatives handling

Since LinkGuard is a rule-based heuristic algorithm, it may cause false positives (i.e., treat non-phishing site as phishing site) and false negatives (i.e., treat phishing site as non-phishing site). In what follows, we show that LinkGuard may result in false positives but is very unlikely to cause harmful false negatives.

For phishing attacks of category 1, we are sure that there is no false positives or false negatives, since the DNS names of the visual and actual links are not the same. It is also easy to observe that LinkGuard handles categories 3 and 4 correctly since the encoded links are first decoded before further analysis.

For category 2, LinkGuard may result in false positives, since using dotted decimal IP addresses instead of domain names may be desirable in some special circumstances (e.g., when the DNS names are still not registered). For category 5, LinkGuard may also result in false positives. For example, we know that both 'www.iee.org' and 'www.ieee.org' are legal Web sites. But these two DNS names have a similarity index of 3/4, hence is very likely to trigger a false positive.

When it is a possible false positive, LinkGuard will return a POSSIBLE_PHISHING. In our implementation (which will be described in the next section), we leverage the user to judge if it is a phishing attack by prompting a dialogue box with detailed information of the hyperlink. The rationale behind this choice is that users generally may have more knowledge of a link than a computer in certain circumstances (e.g., the user may know that the dotted decimal IP address is the address of his friend’s computer and that www.iee.org is a respected site for electrical engineers).

For category 5, LinkGuard may also result in false negatives. False negatives are more harmful than false positives, since attackers in this case will succeed in leading the victim to the phishing sites. For instance, when the sender’s e-mail address and the DNS name in the actual link are the same and the DNS name in the actual link has a very low similarity index with the target site, LinkGuard will return NO_PHISHING. For instance, PatternMatching will treat the below link as NO_PHISHING.

```
<a href="http://fdic-secure.com/application.htm"> Click here </a>
```

with “securehq@fdic-secure.com” as the sender address.

We note that this kind of false negatives is very unlikely to result in information leakage, since the end user is very unlikely to have information the attack interested (since the DNS name in this link is not similar with any legal Web sites).

IV. IMPLEMENTATION AND VERIFICATION OF LINKGUARD

We have implemented the LinkGuard algorithm in Windows XP. It includes two parts: a whook.dll dynamic library and a LinkGuard executive. The structure of the implementation is depicted in Fig. 3.

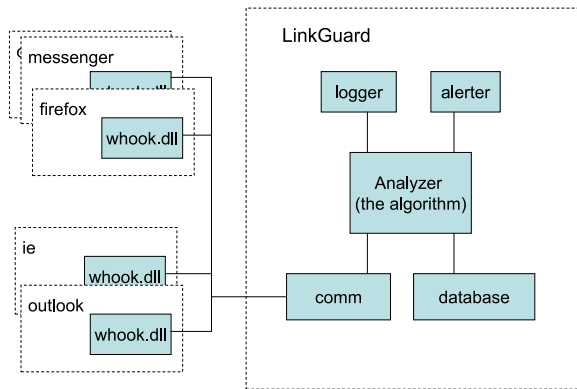


Fig. 3. The structure of the LinkGuard implementation, which consists of a whook.dll and a LinkGuard executive.

whook is a dynamic link library, it is dynamically loaded into the address spaces of the executing processes by the operating system. whook is responsible for collecting data,

such as the called links and visual links, the user input URLs. More specifically, whook.dll is used to: 1) install a BHO (browser helper object) for IE to monitor user input URLs; 2) install an event hook with the *SetWinEventHook* provided by the Windows operating system to collect relevant information; 3) retrieve sender’s e-mail address from Outlook; 4) analyze and filter the received windows and browser events passed by the BHO and the hook, and pass the analyzed data to the LinkGuard executive.

LinkGuard is the key component of the implementation. It is a stand alone windows program with GUI (graphic user interface). It’s composed of 5 parts as illustrated in Fig. 3: Analyzer, Alerter, Logger, Comm, and Database. The functionalities of these 5 parts are given below:

Comm: Communicate with the whook.dll of all of the monitored processes, collect data related to user input from other processes (e.g. IE, outlook, firefox, etc.), and send these data to the Analyzer, it can also send commands (such as block the phishing sites) from the LinkGuard executive to whook.dll. The communication between the LinkGuard process and other processes is realized by the shared memory mechanism provided by the operating system.

Database: Store the whitelist, blacklist, and the user input URLs.

Analyzer: It is the key component of LinkGuard, which implements the LinkGuard algorithm,. It uses data provided by Comm and Database, and sends the results to the Alert and Logger modules.

Alerter: When receiving a warning messages from Analyzer, it shows the related information to alert the users and send back the reactions of the user back to the Analyzer.

Logger: Archive the history information, such as user events, alert information, for future use.

After implemented the LinkGuard system, we have designed experiments to verify the effectiveness of our algorithm. Since we are interested in testing LinkGuard’s ability to detect unknown phishing attacks, we set both whitelist and blacklist to empty in our experiments. Our experiments showed that PhishGuard can detect 195 phishing attacks out of the 203 APWG archives (with detection rate 96%). For the 8 undetected attacks, 4 attacks utilize certain Web site vulnerabilities. Hence the detecting rate is higher than 96% if category 5 is not included. Our experiment also showed that our implementation used by small amount of CPU time and memory space of the system. In a computer with 1.6G Pentium CPU and 512MB memory, our implementation consumes less than 1% CPU time and its memory footprint is less than 7MB.

Our experiment only used the phishing archive provided by APWG as the attack sources. We are planning to use LinkGuard in daily life to further evaluate and validate its effectiveness. Since we believe that a hybrid approach may be more effective for phishing defense, we are also planning to include a mechanism to update the blacklist and whitelist in real-time.

V. CONCLUSION

Phishing has becoming a serious network security problem, causing financial loss of billions of dollars to both consumers and e-commerce companies. And perhaps more fundamentally, phishing has made e-commerce distrusted and less attractive to normal consumers. In this paper, we have studied the characteristics of the hyperlinks that were embedded in phishing e-mails. We then designed an anti-phishing algorithm, LinkGuard, based on the derived characteristics. Since PhishGuard is characteristic based, it can not only detect known attacks, but also is effective to the unknown ones. We have implemented LinkGuard for Windows XP. Our experiment showed that LinkGuard is light-weighted and can detect up to 96% unknown phishing attacks in real-time.

We believe that LinkGuard is not only useful for detecting phishing attacks, but also can shield users from malicious or unsolicited links in Web pages and Instant messages. Our future work includes further extending the LinkGuard algorithm, so that it can handle CSS (cross site scripting) attacks.

REFERENCES

- [1] I. Androutsopoulos, J. Koutsias, K.V. Chandrinos, and C.D. Spyropoulos. An Experimental Comparison of Naive Bayesian and Keyword-Based Anti-Spam Filtering with Encrypted Personal E-mail Message. In *Proc. SIGIR 2000*, 2000.
- [2] The Anti-phishing working group. <http://www.antiphishing.org/>.
- [3] Neil Chou, Robert Ledesma, Yuka Teraguchi, Dan Boneh, and John C. Mitchell. Client-side defense against web-based identity theft. In *Proc. NDSS 2004*, 2004.
- [4] Cynthia Dwork, Andrew Goldberg, and Moni Naor. On Memory-Bound Functions for Fighting Spam. In *Proc. Crypto 2003*, 2003.
- [5] EarthLink. ScamBlocker. <http://www.earthlink.net/software/free/toolbar/>.
- [6] David Geer. Security Technologies Go Phishing. *IEEE Computer*, 38(6):18–21, 2005.
- [7] John Leyden. Trusted search software labels fraud site as 'safe'. http://www.theregister.co.uk/2005/09/27/untrusted_search/.
- [8] Microsoft. Sender ID Framework. <http://www.microsoft.com/mscorp/safety/technologies/senderid/default.aspx>.
- [9] Netcraft. Netcraft toolbar. <http://toolbar.netcraft.com/>.
- [10] PhishGuard.com. Protect Against Internet Phishing Scams. <http://www.phishguard.com/>.
- [11] Jonathan B. Postel. Simple Mail Transfer Protocol. RFC821: <http://www.ietf.org/rfc/rfc0821.txt>.
- [12] Georgina Stanley. Internet Security - Gone phishing. <http://www.cyota.com/news.asp?id=114>.
- [13] Meng Weng Wong. Sender ID SPF. <http://www.openspf.org/whitepaper.pdf>.