

Online Identity Theft: Phishing Technology, Chokepoints and Countermeasures

Aaron Emigh
Radix Labs
ate@radixlabs.com

Revision 1.3, October 3, 2005

Acknowledgments

The author acknowledges sponsorship from the U.S. Department of Homeland Security, Science and Technology Directorate (DHS S&T). Points of view in this document are those of the author and do not necessarily represent the official position of the U.S. Department of Homeland Security or the Science and Technology Directorate. The content of this report was shaped by the members of the Identity Theft Technology Council, a public-private partnership between DHS S&T, SRI International, the Anti-Phishing Working Group (APWG), and private industry. Particular thanks are due to Dan Boneh, Drew Dean, Louie Gasparini, John Mitchell, Robert Rodriguez, Jim Roskind and Don Wilborn for their contributions.

Intended Audience

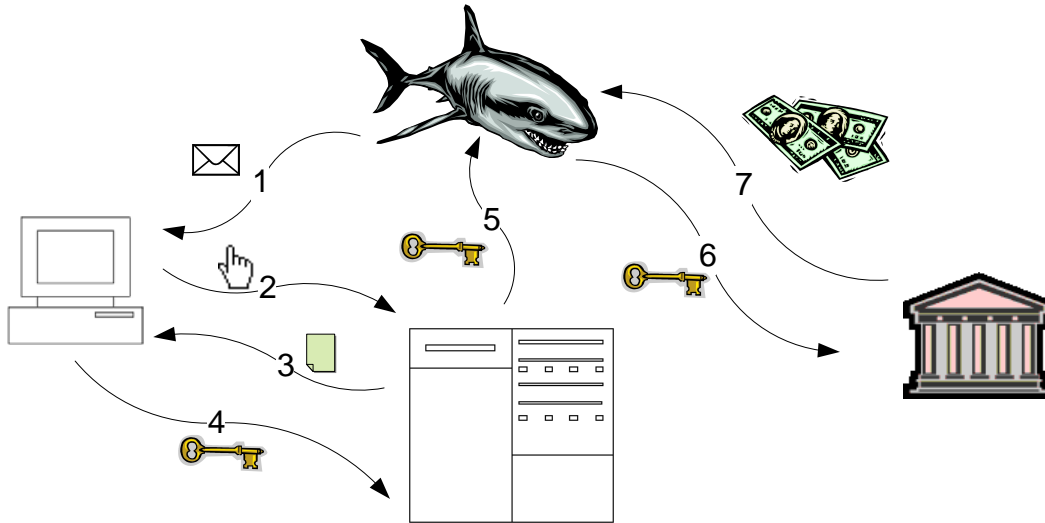
This report is intended for technically sophisticated readers such as security practitioners, executives, researchers, and others who wish to understand methods employed by online identity thieves and countermeasures that can prevent such crimes.

Executive Summary

Phishing is online identity theft in which confidential information is obtained from an individual. Phishing includes deceptive attacks, in which users are tricked by fraudulent messages into giving out information; malware attacks, in which malicious software causes data compromises; and DNS-based attacks, in which the lookup of host names is altered to send users to a fraudulent server.

The Gartner group estimates that the direct phishing-related loss to US banks and credit card issuers in 2003 was \$1.2 billion. Indirect losses are much higher, including customer service expenses, account replacement costs, and higher expenses due to decreased use of online services in the face of widespread fear about the security of online financial transactions. Phishing also causes substantial hardship for victimized consumers, due to the difficulty of repairing credit damaged by fraudulent activity.

This report examines the information flow in phishing attacks of all types. Technologies used by phishers are discussed, in combination with countermeasures that can be applied. The focus is primarily on technology that can be deployed to stop phishing. Both currently available countermeasures and research-stage technologies are discussed.



Steps in a Phishing Attack

All phishing attacks fit into the same general information flow. At each step in the flow, different countermeasures can be applied to stop phishing. The steps are:

0. The phisher prepares for the attack. Step 0 countermeasures include monitoring malicious activity to detect a phishing attack before it begins.
1. A malicious payload arrives through some propagation vector. Step 1 countermeasures involve preventing a phishing message or security exploit from arriving.
2. The user takes an action that makes him or her vulnerable to an information compromise. Step 2 countermeasures involve detecting phishing tactics and rendering phishing messages less deceptive.
3. The user is prompted for confidential information, either by a remote web site or locally by a Web Trojan. Step 3 countermeasures are focused on preventing phishing content from reaching the user.
4. The user compromises confidential information. Step 4 countermeasures concentrate on preventing information from being compromised.
5. The confidential information is transmitted from a phishing server to the phisher. Step 5 countermeasures involve tracking information transmittal.
6. The confidential information is used to impersonate the user. Step 6 countermeasures center on rendering the information useless to a phisher.
7. The phisher engages in fraud using the compromised information. Step 7 countermeasures focus on preventing the phisher from receiving money.

Phishing is complex phenomenon that includes social factors as well as technology. There is no single “silver bullet” that can prevent all phishing. However, properly applied technology can significantly reduce the risk of identity theft. There are many opportunities to apply such technology, including:

- Monitoring potentially malicious activity such as web site usage and domain registrations, detecting a phishing attack before it starts, and interrupting the phisher’s preparations (step 0).
- Authenticating email messages so unauthenticated messages can be discarded (step 1).
- Detecting the unauthorized use of trademarks, logos and other proprietary imagery (step 1).
- Improving the security patching infrastructure to increase resistance to malware (step 1).
- Using personalized information to authenticate an email directly to a user (step 2).
- Detecting a fraudulent web site and alerting the user (step 4).
- Using a mutual authentication protocol (step 4).
- Establishing a trusted path between the user and a web site to ensure that information can be used only by its intended recipient (steps 4 and 6).
- Using two-factor authentication (step 6).
- Forcing passwords to be site-specific (step 6).
- Encoding credentials with restrictions on their validity, using public key cryptography (step 6).

Table of Contents

- Introduction6
- Types of Phishing Attacks.....6
 - Deceptive Phishing.....7
 - Malware-Based Phishing.....8
 - Keyloggers and Screenloggers.....9
 - Session Hijackers9
 - Web Trojans10
 - Hosts File Poisoning10
 - System Reconfiguration Attacks10
 - Data Theft.....10
 - DNS-Based Phishing (“Pharming”).....10
 - Content-Injection Phishing11
 - Man-in-the-Middle Phishing.....11
 - Search Engine Phishing12
- Technology, Chokepoints and Countermeasures13
 - Step 0: Preventing a Phishing Attack Before it Begins14
 - Detecting an Imminent Attack.....14
 - Preparing for an Attack15
 - Step 1: Preventing Delivery of Phishing Payload16
 - Step 1 Countermeasure: Filtering16
 - Step 1 Countermeasure: Email Authentication17
 - Step 1 Countermeasure: Secure Patching19
 - Step 2: Preventing or Disrupting a User Action20
 - Step 2 Countermeasure: Education.....20
 - Step 2 Countermeasure: Use Personalized information21
 - Step 2 Countermeasure: Display Deceptive Content Canonically23
 - Step 2 Countermeasure: Interfere with Navigation26
 - Step 2 Countermeasure: Detect Inconsistent DNS Information27
 - Step 2 Countermeasure: Modify Referenced Images28
 - Steps 2 and 4: Prevent Navigation and Data Compromise28
 - Step 2 and 4 Countermeasure: Increase Inter-Application Data Sharing28

Step 3: Preventing Transmission of the Prompt	29
Step 3 Countermeasure: Filter Out Cross-Site Scripting	29
Step 3 Countermeasure: Disable Injected Scripts	30
Step 4: Preventing Transmission of Confidential Information	31
Step 4 Countermeasure: Anti-Phishing Toolbars.....	32
Step 4 Countermeasure: Data Destination Blacklisting	33
Step 4 Countermeasure: Screen-Based Data Entry	34
Step 4 Countermeasure: Mutual Authentication	34
Steps 4 and 6: Preventing Data Entry and Rendering it Useless	35
Steps 4 and 6 Countermeasure: Trusted Path.....	35
Step 5: Tracing Transmission of Compromised Credentials	38
Step 6: Interfering with the Use of Compromised Information	39
Step 6 Countermeasure: Conventional Two-Factor Authentication	39
Step 6 Countermeasure: Computer-Based Second-Factor Authentication..	40
Step 6 Countermeasure: Password Hashing	41
Step 6 Countermeasure: Transaction Confirmation.....	42
Step 6 Countermeasure: Policy-Based Data	43
Step 7: Interfering with the Financial Benefit	44
Non-Technical Best Practices	45
Conclusions	47
Appendix A: Technology Vendors.....	49
Appendix B: Law Enforcement Resources.....	53
Appendix C: Phishing Bibliography	54
Appendix D: Other Resources	56
Appendix E: The DHS-SRI Identity Theft Technology Council.....	57

Introduction

Phishing is online identity theft in which confidential information is obtained from an individual. It is distinguished from offline identity theft such as card skimming and “dumpster diving,” as well as from large-scale data compromises in which information about many individuals is obtained at once. Phishing includes many different types of attacks, including:

- Deceptive attacks, in which users are tricked by fraudulent messages into giving out information;
- Malware attacks, in which malicious software causes data compromises; and
- DNS-based attacks, in which the lookup of host names is altered to send users to a fraudulent server.

Phishing targets many kinds of confidential information, including user names and passwords, social security numbers, credit card numbers, bank account numbers, and personal information such as birthdates and mothers’ maiden names.

The Gartner group estimates that the direct phishing-related loss to US banks and credit card issuers in 2003 was \$1.2 billion. Indirect losses are much higher, including customer service expenses, account replacement costs, and higher expenses due to decreased use of online services in the face of widespread fear about the security of online financial transactions. Phishing also causes substantial hardship for victimized consumers, due to the difficulty of repairing credit damaged by fraudulent activity.

Both the frequency of phishing attacks and their sophistication is increasing dramatically. Descriptions of recent phishing attacks, and related statistics, may be found at <http://www.antiphishing.org>.

Phishing often spans multiple countries and is commonly perpetrated by organized crime. While legal remedies can and should be pursued by affected institutions, technical measures to prevent phishing are an integral component of any long-term solution.

This report examines technologies employed by phishers and evaluates technical countermeasures, both commercially available and proposed.

Types of Phishing Attacks

Phishing is perpetrated in many different ways. Phishers are technically innovative, and can afford to invest in technology. It is a common misconception that phishers are amateurs. This is not the case for the most dangerous phishing attacks, which are carried out as professional organized crime. As financial institutions have increased their online presence, the economic value of compromising account information has increased dramatically. Criminals such

as phishers can afford an investment in technology commensurate with the illegal benefits gained by their crimes.

Given both the current sophistication and rapid evolution of phishing attacks, a comprehensive catalogue of technologies employed by phishers is not feasible. Several types of attacks are discussed below. The distinctions between attack types are porous, as many phishing attacks are hybrid attacks employing multiple technologies. For example, a deceptive phishing email could direct a user to a site that has been compromised via content injection, which installs malware that poisons the user's hosts file. Subsequent attempts to reach legitimate web sites will be rerouted to phishing sites, where confidential information is compromised using a man-in-the-middle attack.

Deceptive Phishing

While the term "phishing" originated in AOL account theft using instant messaging, the most common vector for deceptive phishing today is email. In a typical scenario, a phisher sends deceptive email, in bulk, with a "call to action" that demands the recipient click on a link. Examples of a "call to action" include:

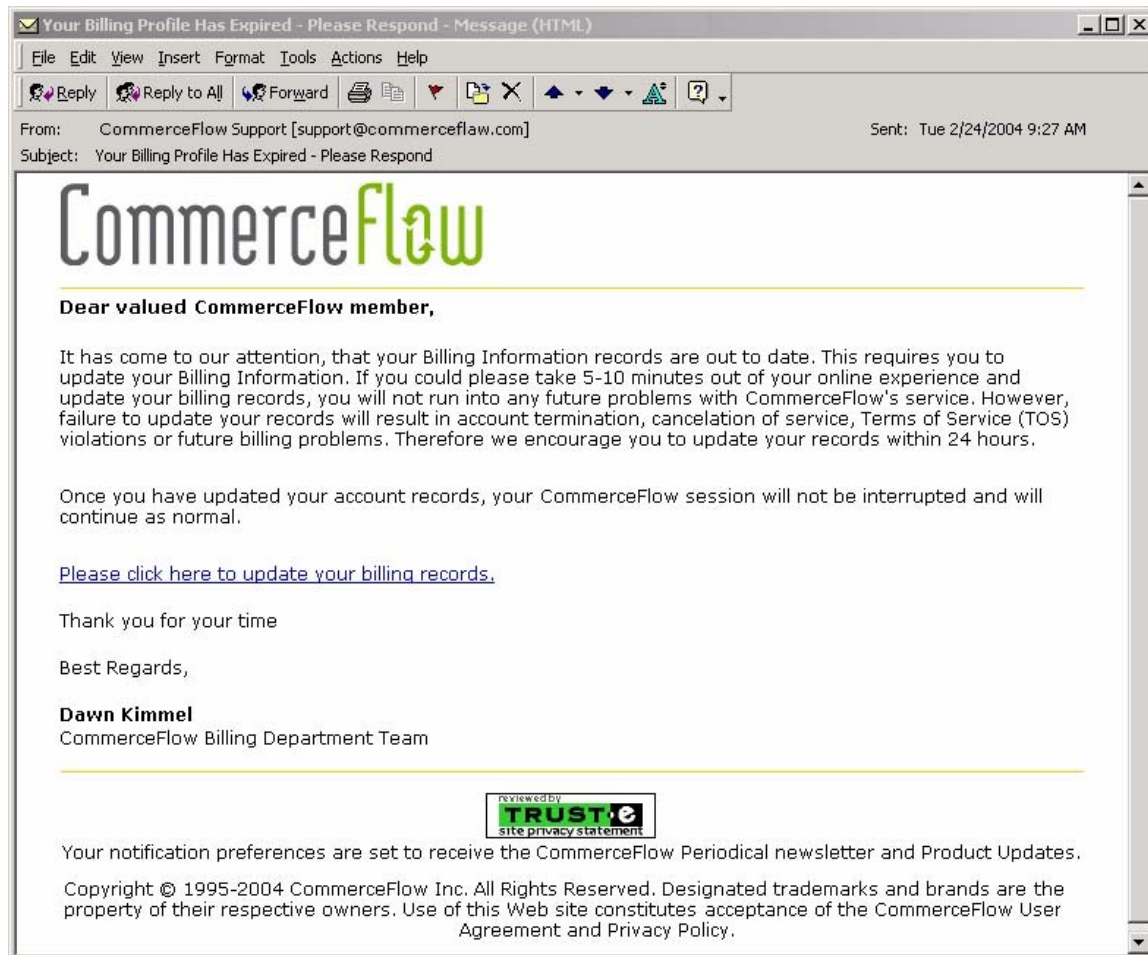
- A statement that there is a problem with the recipient's account at a financial institution or other business. The email asks the recipient to visit a web site to correct the problem, using a deceptive link in the email.
- A statement that the recipient's account is at risk, and offering to enroll the recipient in an anti-fraud program.
- A fictitious invoice for merchandise, often offensive merchandise, that the recipient did not order, with a link to "cancel" the fake order.
- A fraudulent notice of an undesirable change made to the user's account, with a link to "dispute" the unauthorized change.
- A claim that a new service is being rolled out at a financial institution, and offering the recipient, as a current member, a limited-time opportunity to get the service for free.

In each case, the web site to which the user is directed collects the user's confidential information. If a recipient enters confidential information into the fraudulent web site, the phisher can subsequently impersonate the victim to transfer funds from the victim's account, purchase merchandise, take out a second mortgage on the victim's home, file for unemployment benefits in the victim's name, or inflict other damage.

In many cases, the phisher does not directly cause the economic damage, but resells the illicitly obtained information on a secondary market. Criminals participate in a variety of online brokering forums and chat channels where such information is bought and sold.

There are many variations on deception-based phishing schemes. With HTML email readers, it is possible to provide a replica of a login page directly in email, eliminating the need to click on a link and activate the user's web browser.

Sometimes, a numeric IP address is used instead of a host name in a link to a phishing site. In such cases, it is possible to use Javascript to take over the address bar of a browser or otherwise deceive the user into believing he or she is communicating with a legitimate site. A *cousin domain attack* avoids the need for such complexity by using a domain name controlled by a phisher that is deceptively similar to a legitimate domain name, such as www.commerceflow-security.com instead of www.commerceflow.com. Sometimes, an initial deception-based message leads to an installation of malware when a user visits the malicious site.



Typical Deceptive Phishing Message

Malware-Based Phishing

Malware-based phishing refers generally to any type of phishing that involves running malicious software on the user's machine. Malware-based phishing can take many forms. The most prevalent forms are discussed below.

In general, malware is spread either by social engineering or by exploiting a security vulnerability. A typical social engineering attack is to convince a user to open an email attachment or download a file from a web site, often claiming the

attachment has something to do with pornography, salacious celebrity photos or gossip. Some downloadable software can also contain malware. Malware is also spread by security exploits either by propagating a worm or virus that takes advantage of a security vulnerability to install the malware, or by making the malware available on a web site that exploits a security vulnerability. Traffic may be driven to a malicious web site via social engineering such as spam messages promising some appealing content at the site, or by injecting malicious content into a legitimate web site by exploiting a security weakness such as a cross-site scripting vulnerability on the site.

Keyloggers and Screenloggers

Keyloggers are programs that install themselves either into a web browser or as a device driver, which monitor data being input and send relevant data to a phishing server. Keyloggers use a number of different technologies, and may be implemented in many ways, including:

- A browser helper object that detects changes to the URL and logs information when a URL is at a designated credential collection site;
- A device driver that monitors keyboard and mouse inputs in conjunction with monitoring the user's activities; and
- A *screenlogger* that monitors both the user's inputs and the display to thwart alternate on-screen input security measures.

Keyloggers may collect credentials for a wide variety of sites. Keyloggers are often packaged to monitor the user's location and only transmit credentials for particular sites. Often, hundreds of such sites are targeted, including financial institutions, information portals, and corporate VPNs. Various secondary damage can be caused after a keylogger compromise. In one real-world example, the inclusion of a credit reporting agency in a keylogger spread via pornography spam led to the compromise of over 50 accounts with access to the agency, which in turn were ultimately used to compromise over 310,000 sets of personal information from the credit reporting agency's database.

Session Hijackers

Session hijacking refers to an attack in which a user's activities are monitored, typically by a malicious browser component. When the user logs into his or her account, or initiates a transaction, the malicious software "hijacks" the session to perform malicious actions once the user has legitimately established his or her credentials.

Session hijacking can be performed on a user's local computer by malware, or can also be performed remotely as part of a man-in-the-middle attack, which will be discussed later. When performed locally by malware, session hijacking can look to the targeted site exactly like a legitimate user interaction, being initiated from the user's home computer.

Web Trojans

Web Trojans are malicious programs that pop up over login screens to collect credentials. The user believes that he or she is entering information on a web site, while in fact the information is being entered locally, then transmitted to the phisher for misuse.

Hosts File Poisoning

If a user types www.company.com into his or her URL bar, or uses a bookmark, the user's computer needs to translate that address into a numeric address before visiting the site. Many operating systems, such as Windows, have a shortcut "hosts" file for looking up host names before a DNS (Domain Name System) lookup is performed. If this file is modified, then www.company.com can be made to refer to a malicious address. When the user goes there, he or she will see a legitimate-looking site and enter confidential information, which actually goes to the phisher instead of the intended legitimate site.

System Reconfiguration Attacks

System reconfiguration attacks modify settings on a user's computer to cause information to be compromised.

One type of system reconfiguration attack is to modify a user's DNS servers, so faulty DNS information can be provided to users as described below.

Another type of system reconfiguration attack is to install a web proxy, through which the user's traffic will be passed. This is a form of a man-in-the-middle attack, which is discussed separately.

Data Theft

Once malicious code is running on a user's computer, it can directly steal confidential information stored on the computer. Such information can include passwords, activation keys to software, sensitive email, and any other data that is stored on a victim's computer. By automatically filtering data looking for information that fits patterns such as a social security number, a great deal of sensitive information can be obtained. Data theft is also widely used for phishing attacks aimed at corporate espionage, based on the fact that personal computers often contain the same confidential information that is also stored on better-protected enterprise computers. In addition to espionage for hire, confidential memos or design documents can be publicly leaked, causing economic damage or embarrassment.

DNS-Based Phishing ("Pharming")

DNS-based phishing is used here to refer generally to any form of phishing that interferes with the integrity of the lookup process for a domain name. This includes hosts file poisoning, even though the hosts file is not properly part of the Domain Name System. Hosts file poisoning is discussed in the malware section since it involves changing a file on the user's computer.

Another form of DNS-based phishing involves polluting the user's DNS cache with incorrect information that will be used to direct the user to an incorrect location. If the user has a misconfigured DNS cache, this can be done directly. It can also be done with a system reconfiguration attack that changes the user's DNS server to a malicious server, by hacking a legitimate DNS server, or by polluting the cache of a misconfigured legitimate DNS server.

Content-Injection Phishing

Content-injection phishing refers to inserting malicious content into a legitimate site. The malicious content can redirect to other sites, install malware on a user's computer, or insert a frame of content that will redirect data to a phishing server.

There are three primary types of content-injection phishing, with many variations of each:


- Hackers can compromise a server through a security vulnerability and replace or augment the legitimate content with malicious content.
- Malicious content can be inserted into a site through a cross-site scripting vulnerability. A cross-site scripting vulnerability is a programming flaw involving content coming from an external source, such as a blog, a user review of a product on an e-commerce site, an auction, a message in a discussion board, a search term or a web-based email. Such externally supplied content can be a malicious script or other content that is not properly filtered out by software on the site's server, and runs in the web browser of a visitor to the site.
- Malicious actions can be performed on a site through a SQL injection vulnerability. This is a way to cause database commands to be executed on a remote server that can cause information leakage. Like cross-site scripting vulnerabilities, SQL injection vulnerabilities are a result of improper filtering.

Cross-site scripting and SQL injection are propagated through two different primary vectors. In one vector, malicious content is injected into data stored on a legitimate web server, such as an auction listing, product review or web-based email. In the other vector, malicious content is embedded into a URL that the user visits when he or she clicks on a link. This is commonly a URL that will be displayed on screen or used as part of a database query, such as an argument to a search function.

Man-in-the-Middle Phishing

A man-in-the-middle attack is a form of phishing in which the phisher positions himself between the user and the legitimate site. Messages intended for the legitimate site are passed to the phisher instead, who saves valuable information, passes the messages to the legitimate site, and forwards the responses back to the user. Man-in-the-middle attacks can also be used for session hijacking, with or without storing any compromised credentials. Man-in-the-middle attacks are

difficult for a user to detect, because the site will work properly and there may be no external indication that anything is wrong.



Man-in-the-Middle Attack

Man-in-the-middle attacks may be performed using many different types of phishing. Some forms of phishing, such as proxy attacks, are inherently man-in-the-middle attacks. However, man-in-the-middle attacks may be used with many other types of phishing, including DNS-based phishing and deception-based phishing.

Normally, SSL web traffic will not be vulnerable to a man in the middle. The handshake used by SSL ensures that the session is established with the party named in the server's certificate, and that an attacker cannot obtain the session key; and SSL traffic is encrypted using the session key so it cannot be decoded by an eavesdropper. Proxies have a provision for tunneling such encrypted traffic. However, a malware-based attack can modify a system configuration to install a new trusted certificate authority, in which case such a man in the middle can create its own certificates for any SSL-protected site, decrypt the traffic and extract confidential information, and re-encrypt the traffic to communicate with the other side. In practice, man-in-the-middle attacks simply do not use SSL, since users do not generally check for its presence.

Man-in-the-middle attacks can also compromise authentication credentials, such as one-time or time-varying passcodes generated by hardware devices. Such stolen credentials can be used by the phisher for authentication as long as they remain valid.

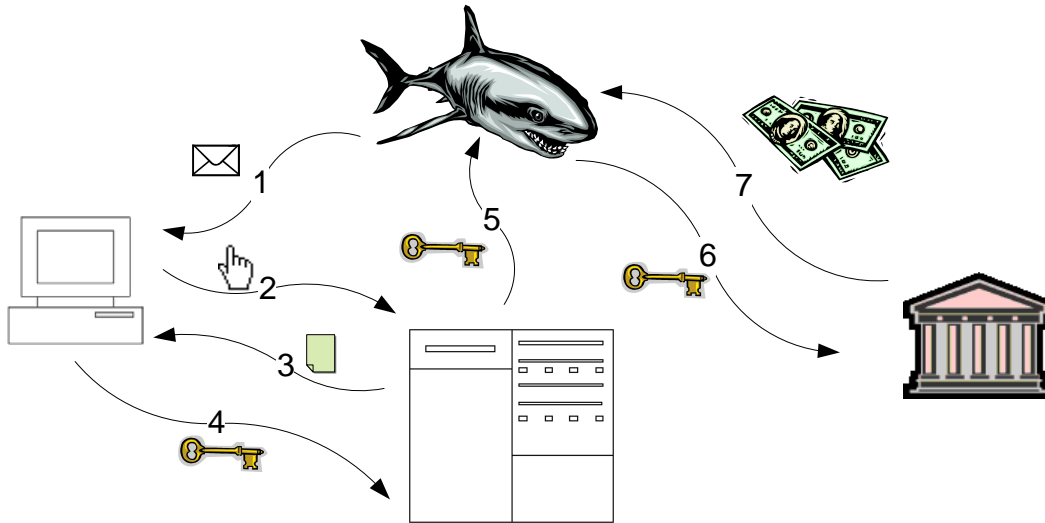
Search Engine Phishing

Another approach taken by phishers is to create web pages for fake products, get the pages indexed by search engines, and wait for users to enter their confidential information as part of an order, sign-up, or balance transfer. Such pages typically offer products at a price slightly too good to be true.

Scams involving fraudulent banks have been particularly successful. A phisher creates a page advertising an interest rate slightly higher than any real bank. Victims find the online bank via a search engine, and enter their bank account credentials for a "balance transfer" to the new "account." Greed is a powerful motivator that can cloud judgment. Some victims even provided their bank account numbers to "Flintstone National Bank," of "Bedrock, Colorado."

Technology, Chokepoints and Countermeasures

Technology may be applied to stop a phishing attack at multiple stages. Technology countermeasures are discussed with reference to the steps in the following information flow of a phishing attack.



Steps in a Phishing Attack

Step by step according to the figure above, the fundamental flow of information in a phishing attack is:

0. The phisher prepares for the attack. For certain types of attacks, such as deceptive attacks using cousin domains, a domain must be registered. Phishing servers are established, either owned by the phisher or (more often) computers that have been compromised by hacking or malware. Phishing servers are configured to receive information, whether from the user in a web-based interface or from malware on victims' computers.
1. A malicious payload arrives through some propagation vector. In a deception-based phishing attack, the payload is typically a deceptive email. In the case of a malware or system reconfiguration attack, the payload is malicious code that arrives as an attachment to an email, an unintended component of downloaded software, or an exploit of a security vulnerability. In the case of a DNS poisoning attack, the payload is false IP address information. In the case of search engine phishing, the payload is a search result referencing a fraudulent site. In the case of a cross-site scripting attack, the payload is malicious code that is stored on a legitimate server or embedded in a URL in an email, depending on the attack details.
2. The user takes an action that makes him or her vulnerable to an information compromise. In a deception-based phishing attack, the user clicks on a link. In a keylogger attack, the user goes to a legitimate web site. In a host name lookup attack, the user goes to a legitimately-named site that is diverted to a fraudulent site.

3. The user is prompted for confidential information, either by a remote web site or locally by a Web Trojan. A remote web site sending the prompt can be a legitimate site (in the case of a keylogger attack), or a malicious site (in the case of a deception-based attack or a DNS attack), or a legitimate web site providing malicious code (in the case of a content-injection attack).
4. The user compromises confidential information such as a credential, either by providing it to a malicious server, to malicious software running locally, or to software that is eavesdropping on a legitimate interaction.
5. The confidential information is transmitted to the phisher. Depending on the nature of the attack, this information can be sent by a malicious or compromised server, or in the case of locally running malware such as a keylogger or Web Trojan, the information can be sent by the victim's PC.
6. The confidential information is used to impersonate the user.
7. A fraudulent party obtains illicit monetary gain, or otherwise engages in fraud using the confidential information.

Each step in the phishing information flow is examined. At each step, technology countermeasures are evaluated that can be applied to stop a phishing attack at that juncture.

Step 0: Preventing a Phishing Attack Before it Begins

In some cases, it may be possible to detect a phishing attack before it occurs. A company can also prepare for a phishing attack in the absence of a crisis, to improve responsiveness and mitigate losses.

Detecting an Imminent Attack

To carry out some kinds of phishing attacks, such as deceptive attacks using cousin domains, a phisher must set up a domain to receive phishing data. Pre-emptive domain registrations targeting likely spoof domain names may reduce the availability of the most deceptively named domains.

Since there may be millions of possible spoofing domains, it is not generally practical to register all possible official-looking domains. Some companies offer a registration monitoring service that will detect registration of a potential spoof domain and monitor any site activity while pursuing action against the registrant.

Proposals have been made to institute a "holding period" for new domain registrations, during which trademark holders could object to a new registration before it was granted. This might help with the problem of cousin domains, but would not address the ability of phishers to impersonate sites.

Setting up a phishing server often involves saving a copy of the legitimate site that is being impersonated. It is sometimes possible to analyze access patterns in web logs on a legitimate site and detect phishers' downloading activities. While pages on a public web site cannot ultimately be kept from phishers, this can provide lead time on responding to an attack, and an early analysis based on

the IP addresses being used can sometimes accelerate an investigation once an attack is underway.

Some services attempt to search the web and identify new phishing sites before they go “live.” Such services can often result in shutting down a phishing site before it is active. In many cases, however, phishing sites may not be accessible to search spiders, and do not need to be active for long, as most of the revenues are gained in the earliest period of operation. The average phishing site stays active no more than two days, often only a matter of hours, yet that is sufficient to collect substantial revenues.

Phishers have deployed a variety of technologies to keep phishing servers online for longer periods of time. For example, phishing using a domain that a phisher owns can be directed to arbitrary IP addresses by updating information on the DNS servers for the phishing domain. Phishers have set up custom DNS servers and rotated between them, providing IP addresses in a round robin fashion for many compromised machines. Whenever a phishing server is taken down, it is removed from the rotation and another compromised machine is added. Whenever a DNS server is taken down, the registration information is modified to replace it with another one. This has the effect of requiring a takedown through the domain registrar, which can be a more cumbersome and time-consuming effort than taking down a machine through an ISP. Some phishers also set up port redirectors on compromised machines to which victims are sent, to function as load balancers and allow replacement of phishing servers as they are taken down.

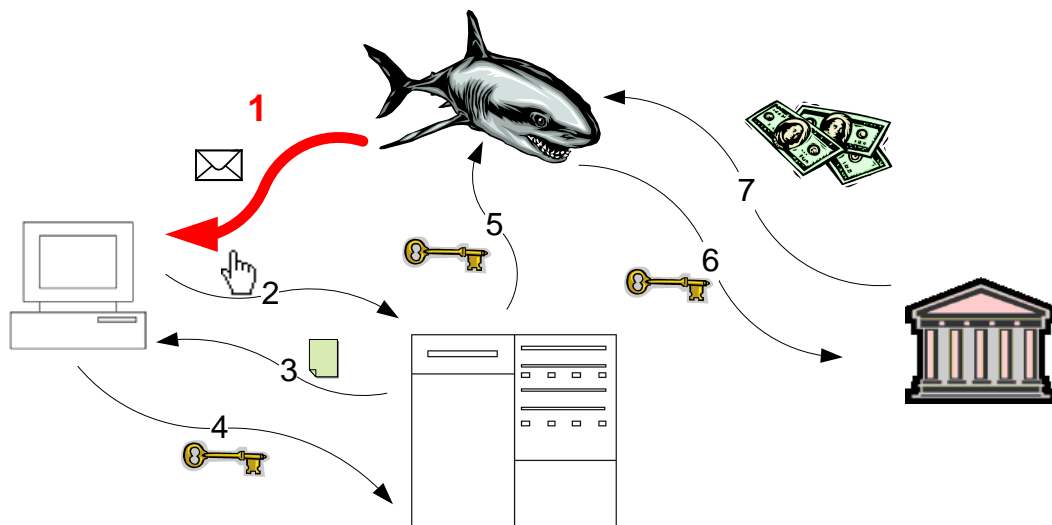
Preparing for an Attack

Before an attack occurs, an organization that is a likely phishing target can prepare for an attack. Such preparation can dramatically improve the organization’s responsiveness to the attack and reduce losses substantially. Such preparation includes:

- Providing a spoof-reporting email address that customers may send spoof emails to. This may both provide feedback to customers on whether communications are legitimate, and provide warning that an attack is underway.
- Monitoring “bounced” email messages. Many phishers email bulk lists that include nonexistent email addresses, using return addresses belonging to the targeted institution. A spate of bounced emails can indicate that a phishing attack is underway.
- Monitoring call volumes and the nature of questions to customer service. A spike in certain types of inquiries, such as a password having been changed, can indicate a phishing attack.
- Monitoring account activity for anomalous activity such as unusual volumes of logins, password modification, transfers, withdrawals, etc.

- Monitoring the use of images containing an institution’s corporate logos and artwork. Phishers will often use the target corporation to host artwork that is used to deceive customers. This may be detected by a web server via a blank or anomalous “referrer” for the image.
- Establishing “honeypots” and monitoring for email purporting to be from the institution.

There are contractors that can perform many of these services. Knowing when an attack is underway can be valuable, in that it may permit a targeted institution to institute procedural countermeasures, initiate an investigation with law enforcement, and staff up to respond to the attack in a timely manner.



Phishing Information Flow, Step 1

Step 1: Preventing Delivery of Phishing Payload

Once a phishing attack is underway, the first opportunity to prevent a phishing attack is to prevent a phishing payload, such as an email or security exploit, from ever reaching users. This represents a disruption of step 1 of the phishing information flow.

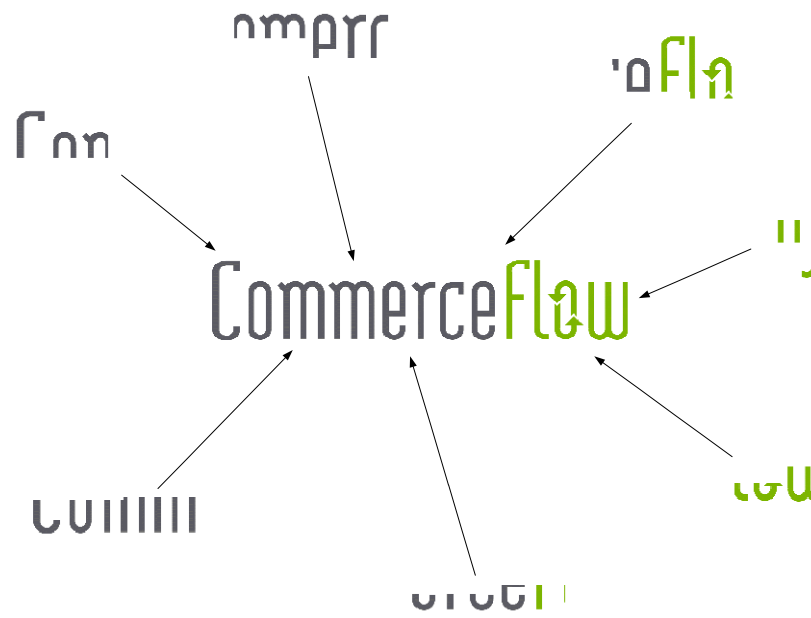
Step 1 Countermeasure: Filtering

Email filters intended to combat spam are often effective in combating phishing as well. Signature-based anti-spam filters may be configured to identify specific known phishing messages and prevent them from reaching users. Statistical or heuristic anti-spam filters may be partially effective against phishing, but to the extent that a phishing message resembles a legitimate message, there is a danger of erroneously blocking legitimate email if the filter is configured to be sufficiently sensitive to identify phishing email.

Effective deception-based phishing emails and web sites must present a visual appearance consistent with the institutions that they are mimicking. Color schemes and imagery mimic the targeted institution. An important aspect of this

is the use of a corporate logo; this dramatically increases the deceptiveness of a phishing email.

One possible countermeasure is to detect unauthorized logos in emails. There are many countermeasures that phishers may employ against a simple image comparison, including displaying many tiled smaller images as a single larger image, and stacking up transparent images to create a composite image.



Composite Logotype Rendering

To avoid such workarounds from phishers, imagery should be fully rendered before analysis. An area of future research is how to recognize potentially modified trademarks or other registered imagery within a larger image such as a fully rendered email. A similar approach may be fruitful when applied to web sites, when a user has clicked on a link.

Step 1 Countermeasure: Email Authentication

Phishing emails typically claim to come from a trusted source. There are two primary ways in which this is accomplished:

- Forging a return address;
- Registering a cousin domain (e.g. “commerceflow-security.com” to spoof a company whose real domain is “commerceflow.com”) and sending email from that domain name.

Message authentication technologies have considerable promise for anti-phishing applications. In general, message authentication provides an assurance that an email was really sent by the party named as the sender. Once widely deployed, email authentication has the potential to prevent forgery of a return address, and force a phisher to either reveal a suspicious-looking return address, or register an official-looking domain name. The advantages of this are

that the return address may be less deceptive than a forged address, a domain registration may be detected in advance of a phishing attack, and a phisher may be traced through the domain registration.

There are many proposals for email authentication technologies. Sender-ID and SPF prevent return address forgery by checking DNS records to determine whether the IP address of a transmitting mail transfer agent (MTA) is authorized to send a message from the sender's domain. Domain Keys and Internet Identified Email provide similar authentication, using a domain-level cryptographic signature that can be verified through DNS records. MTA authorization approaches have the advantage of ease of implementation, while cryptographic approaches offer end-to-end authentication. Sender-ID and SPF are now IETF Experimental Standards, while the MASS working group is working to merge Domain Keys and Internet Identified Email. Other proposals have been made for repudiable cryptographically signed emails, and for authority-based email authentication in which an authentication token certified by an authority can be interpreted by a recipient.

Another approach to email authentication is for a sender to provide a proof of authorization to send an email to the recipient. Such schemes include the automatic generation and use of sender-specific or policy-based email addresses, and the use of a token or certificate issued by the message recipient, granting the sender permission to send. Such approaches require either additional user interfaces (in the case of generation of sender-specific email addresses) or infrastructure (in the case of token generation and/or certificate signing and distribution).

Some form of lightweight message authentication may be very valuable in the future to combat phishing. For the potential value to be realized, email authentication technology must become sufficiently widespread that non-authenticated messages can be summarily deleted or otherwise treated prejudicially, and security issues surrounding the use of mail forwarders in MTA authorization schemes such as Sender-ID need to be resolved.

Cryptographic signing of email (e.g. S/MIME signing) is a positive incremental step in the short run, and an effective measure if it becomes widely deployed in the long run. Signing may be performed either at the client or at the gateway. However, current email clients simply display an indication of whether an email is signed. A typical user is unlikely to notice that an email is unsigned and avoid a phishing attack. Signing could be more effective if the functionality of unsigned emails were reduced, such as by warning when a user attempts to follow a link in an unsigned email. However, this would place a burden on unsigned messages, which today constitute the vast majority of email messages. If critical mass builds up for signed emails, such measures may become feasible.

Step 1 Countermeasure: Secure Patching

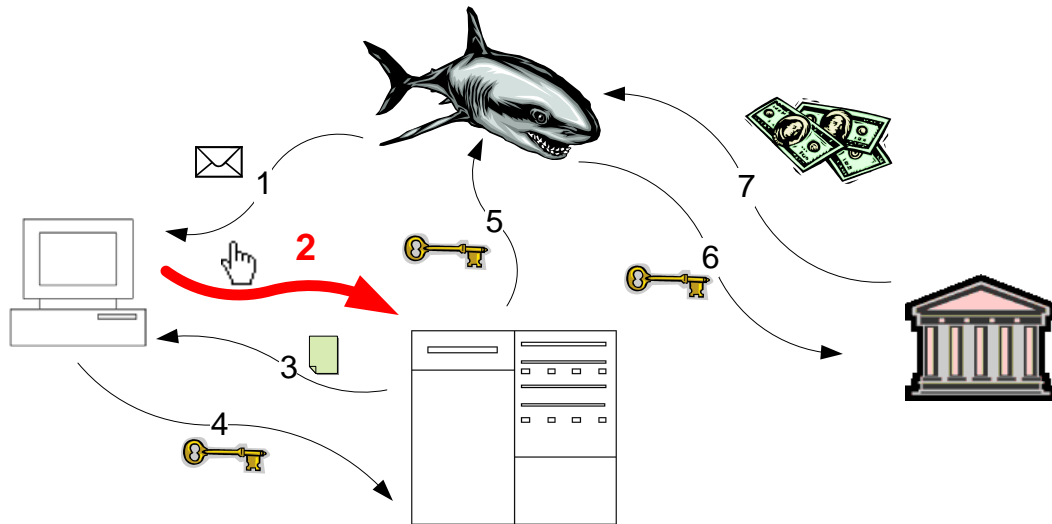
Phishing attacks that involve malware are often installed via an exploit of a security vulnerability. A user running an unpatched operating system or browser runs the risk of becoming infected with malware by browsing or even by simply being connected to the internet.

Almost all exploits target known vulnerabilities. “Zero-day” attacks targeting previously unknown vulnerability are very rare in practice. Therefore, a fully patched computer behind a firewall is the best defense against exploit-based malware installation.

Patches can be large and typically take a long time to be distributed across a worldwide customer base, and users and IT departments often do not apply patches promptly. Studies have shown that it is often wise to wait before applying a patch, to allow time for corrections to an initially buggy patch that could destabilize a computer.

However, announcement and distribution of a patch provides information to criminals about the security vulnerability that is being patched. Even if the description is vague, a patch can be disassembled and compared to the code that it replaces. Once a new exploit is known, a malware exploit can be quickly crafted using pre-built components. It currently takes less than three days – sometimes only a matter of hours – between the time a patch is released and the time a malicious exploit appears. After this short period of time, most computers are still vulnerable to infection.

One promising proposal for rapid distribution and application of patches, without leaking vulnerability information, is to distribute focused security patches for specific vulnerabilities encrypted using a separate symmetric key for each patch. The key is kept secret by the vendor. The patches cannot be applied while encrypted, but they can be distributed to all vulnerable computers without leaking information about the vulnerability to criminals. When an actual exploit of a vulnerability repaired by a patch is detected, the decryption key for that particular patch can be quickly distributed to all computers on the internet for automatic installation of the patch. The exploit could be detected by a version of the patch running on honeypot machines that detects an attempt to exploit the vulnerability that the patch fixes.



Phishing Information Flow, Step 2

Step 2: Preventing or Disrupting a User Action

Step 2 of the phishing information flow involves a user action that takes the user to a location where his or her confidential information may be compromised. Several countermeasures can disrupt this process.

Step 2 Countermeasure: Education

The most widely deployed step 2 countermeasure is to “educate” the user base by instructing users not to click on links in an email, to ensure that SSL is being used, to verify that the domain name is correct before giving out information, and similar practices.

Such education has not been effective: response rates to phishing messages are comparable to response rates to legitimate commercial email. There are at least four likely reasons why this form of education has not proven effective:

- The information normally presented to a user – including the origin of an email, the location of a page, the presence of SSL, etc. – can be spoofed. Therefore, a user, however well-educated, cannot reasonably be relied on to discern between a legitimate message and a phishing attack.
- Actions such as ensuring SSL is being used and checking the domain name are not directly related to a user’s normal interactions with a site, which have been found to make them very likely to be skipped.
- Financial institutions have widely deviated from the guidelines they have disseminated for distinguishing phishing messages from legitimate communications, undermining the educational messages they have distributed. In particular, many financial institutions use unexpected domain names similar to the names a phisher would use, do not use SSL in a user-verifiable way on a login page, include clickable links in email communications, and so on.

- Users are accustomed to glitches and malfunctions, and often are not sure how to interpret phishing-related behavior. Users often rationalize phishing indicators as being due to software bugs or other errors.

Following consistent practices that differ from phishers is likely the most effective way to educate customers, in that customers will become acclimated to a particular mode of interaction with legitimate sites and more suspicious of sites that deviate from such practices. Financial institutions can foster such education by adhering to practices such as:

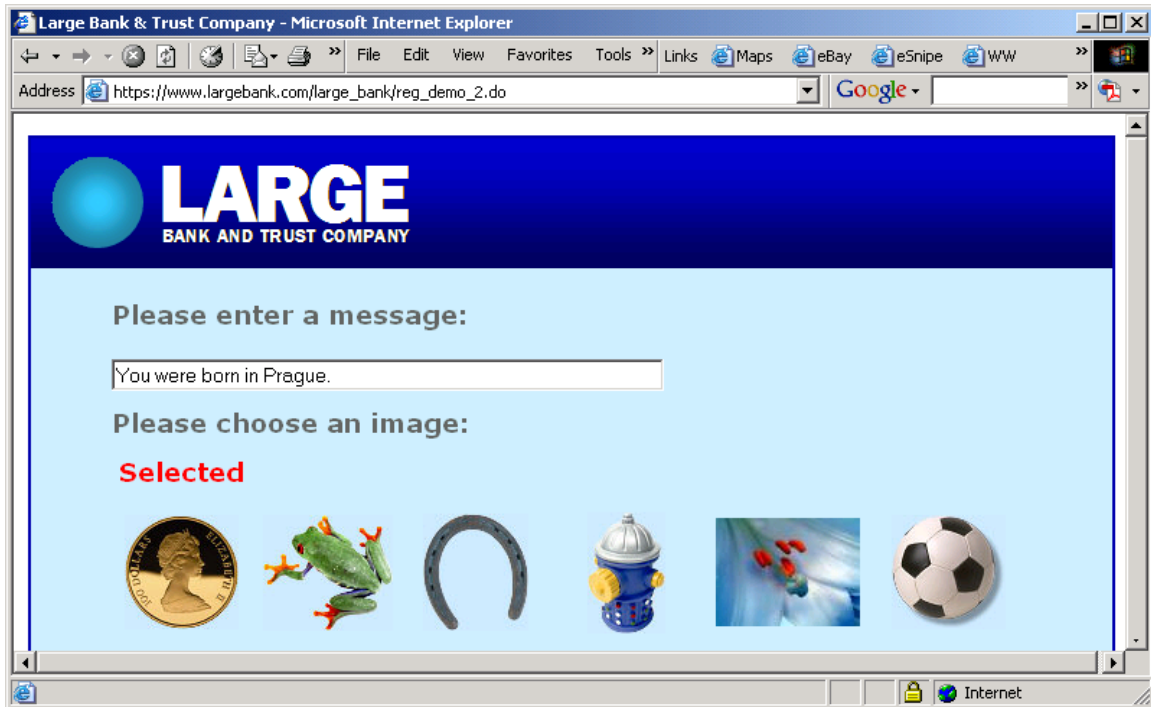
- Not telling customers they will never use clickable links, when in fact such links are a valuable form of marketing;
- Never using a “call to action” in email that warns of a negative consequence for failing to follow a link;
- Using email authentication technology;
- Using honestly named links, if links are used (e.g. not using deceptively named links);
- Always using the expected domain name for logging in; and
- Always using SSL on the login page, and all other pages.

Step 2 Countermeasure: Use Personalized information

A simple way to reduce the deceptiveness of phishing messages is to include personalized information with all legitimate communications. For example, if every email from commerceflow.com begins with the user’s name, and every email from commerceflow.com educates the user about this practice, then an email that does not include a user’s name is suspect. While implementing this practice can be complex due to the difficulty of coordinating multiple business units, affiliate marketing programs and the widespread practice of outsourcing email to external services, it is an effective measure. Since the information may be shared with partners and is generally sent over insecure channels, any personalized information used should not be sensitive.

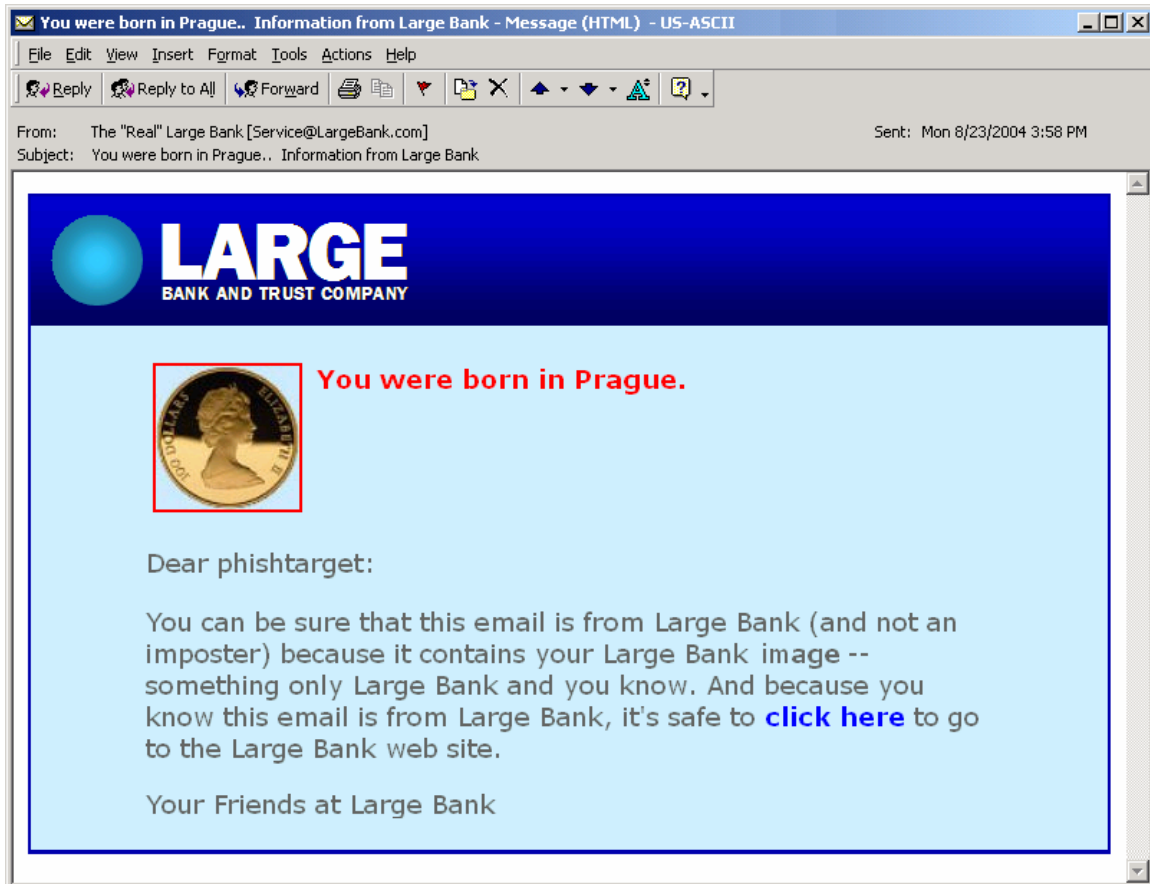
Beyond static identifying information, more sophisticated personalized information may be included, such as text that a user has requested to be used. This permits a user to easily verify that the desired information is included.

Personalized imagery may also be used to transmit messages. For example, when a user creates or updates account information, he or she may be allowed (or required) to enter textual and/or graphical information that will be used in subsequent personalized information. In this example, a customer of the Large Bank and Trust Company has typed in the personalized text “You were born in Prague” and selected or uploaded a picture of a Canadian penny.



Personalized information: Sign-Up

A subsequent email from Large Bank and Trust Company will include this personalized information, e.g.



Personalized information: Email

Since phishers will not know what personalized information a user has selected, they will not be able to forge deceptive emails.

A similar approach can be used for web sites after a user enters a user name, but before entering a password. However, a web site should first authenticate the user by other means. To avoid a man-in-the-middle attack, additional authentication, such as two-factor authentication, should be used to ensure that the user and computer are legitimate before displaying personalized information. When the user is confirmed, personalized text and/or imagery is displayed, and the user enters password information only after verifying that the personalized information is correct.

This type of approach does rely on some user education, but unlike admonitions to check a lock icon, distrust an unsigned email, or type in a URL, there are structural differences in the interaction between a user and a message or site. These structural differences may make a user more likely to discern differences between a phishing attack and a legitimate interaction.

Step 2 Countermeasure: Display Deceptive Content Canonically

A deception-based phishing email typically requires a user to click on a link to go to a web site. The phisher's web site usually does not have a legitimate name,

so the actual destination of the link is often disguised. (Exceptions to this rule include attacks using cousin domains, phishing sites reached through compromises in DNS name resolution, and homograph attacks using Internationalized Domain Names.)

Presently, links may be displayed however the author of the content specifies. This makes it easy to create deceptive links in phishing email. Phishers employ many technologies to obscure the true destination of a link. Examples include:

Misleadingly named links – A link may display as

<http://security.commerceflow.com> but actually lead to <http://phisher.com>.

Cloaked links – URLs can incorporate a user name and password. This can be used to “cloak” the actual destination of a link. For example, the URL

<http://security.commerceflow.com@phisher.com> actually leads to <http://phisher.com>.

Redirected links – “Redirects” that translate a reference to one URL into another URL are commonly used in web programming. If a careless programmer at a targeted institution leaves an “open redirect” accessible that can be used to redirect to an arbitrary location, this can be used by phishers to provide a legitimate-looking URL that will redirect to a phishing site.

Obfuscated links – URLs can contain encoded characters that hide the meaning of the URL. This is commonly used in combination with other types of links, for example to obscure the target of a cloaked or redirected link.

Programmatically obscured links – If scripts are allowed to run, Javascript can change the status text when the user mouses over a link to determine its destination.

Map links – A link can be contained within an HTML “image map” that refers to a legitimate-looking URL. However, the actual location to which a click within the image map directs the browser will not be displayed to the user.

Homograph URLs – A URL in a link can use an IDN (Internationalized Domain Names) homograph, a character that is displayed the same as a regular character but is actually different, typically a character from a different alphabet such as Cyrillic. This is presently a problem mostly for browsers with non-standard configurations.

One possible countermeasure for implementation in an email client or browser is to render potentially deceptive content in a predictable way that clearly identifies it to the user as suspicious. For example, consider the following HTML fragment and a typical rendering of it:

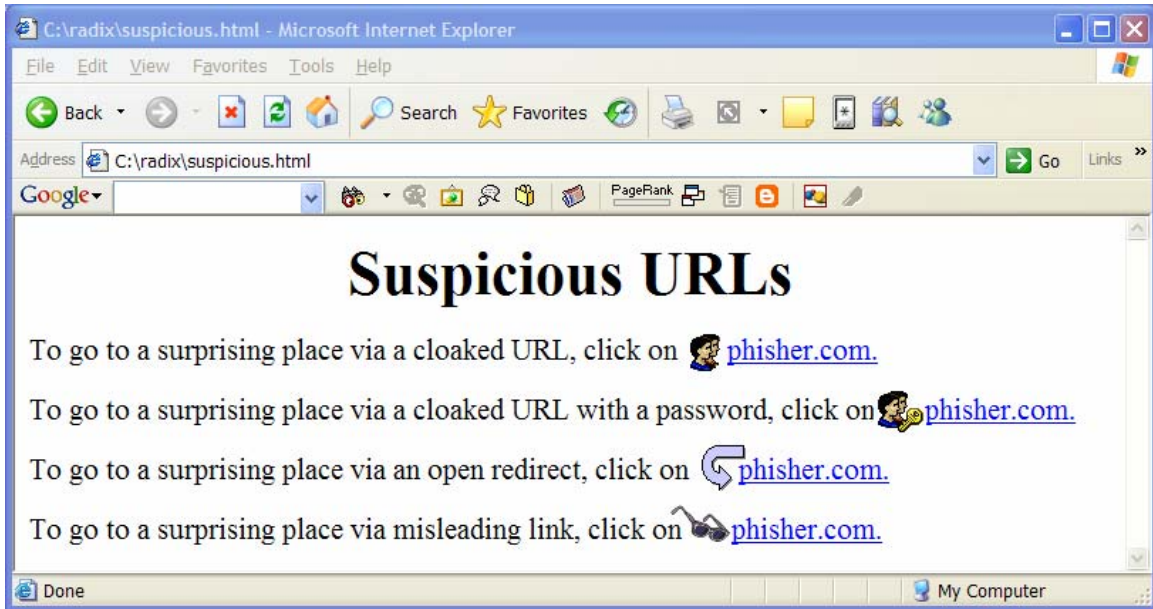
<CENTER><H1>Suspicious URLs</H1></center>
<P>To go to a surprising place via a cloaked URL, click on
this link.
<P>To go to a surprising place via a cloaked URL with a password, click on
this
link.
<P>To go to a surprising place via an open redirect, click on
this link.
<P>To go to a surprising place via misleading link, click on
http://security.commerceflow.com.

HTML Content with Deceptive Links



Rendered HTML Content with Deceptive Links

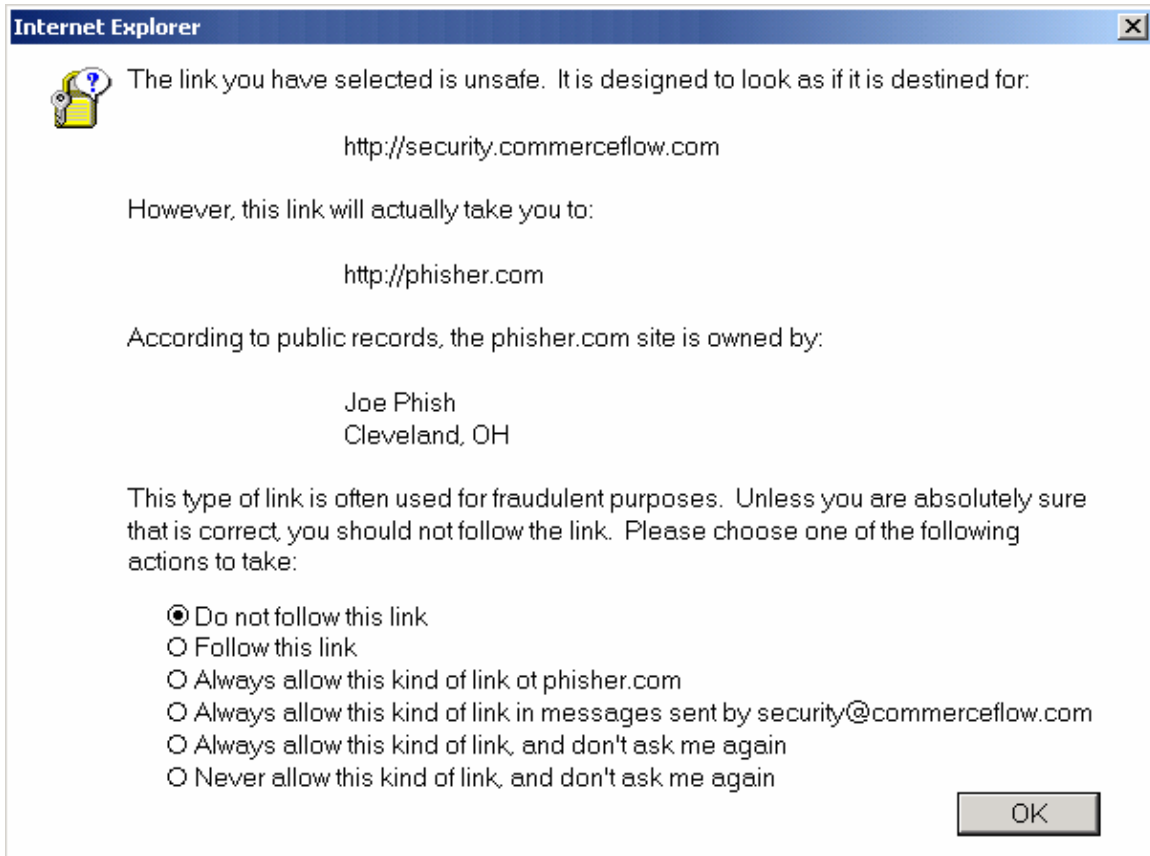
Even looking at the URL in the status bar before clicking, the user may not understand the actual destination of the link he or she is clicking on. This is especially true when link obfuscation is used. An email client or browser extension to iconically show the destination of potentially confusing URLs could clarify the situation for a user, especially if combined with countermeasures for status bar spoofing (for example, always showing the important parts of a URL and not allowing scripts to modify the status bar when a URL is being shown). The page above might be rendered more informatively as:



Rendered HTML Content with Deceptive Links, Displayed Canonically

Step 2 Countermeasure: Interfere with Navigation

When a user clicks on a link that is suspicious, such as a cloaked, obfuscated, mapped, or misleadingly named link, a warning message can be presented advising the user of the potential hazards of traversing the link. Information should be presented in a straightforward way, but need not be simplistic. To help the user make an informed decision, data from sources such as reverse DNS and WHOIS lookups can be usefully included:



Unsafe Link Traversal Warning Message

An informative warning has the benefit of allowing legitimate links even if of a suspicious nature, while providing a risk assessment with the information a user needs to determine an appropriate action.

Studies have shown that such information is more reliably evaluated by a user if it is part of the “critical action sequence” that a user must perform in order to achieve an action. Therefore, an interaction that requires a user to select the intended destination from among several destinations may be more effective.

Step 2 Countermeasure: Detect Inconsistent DNS Information

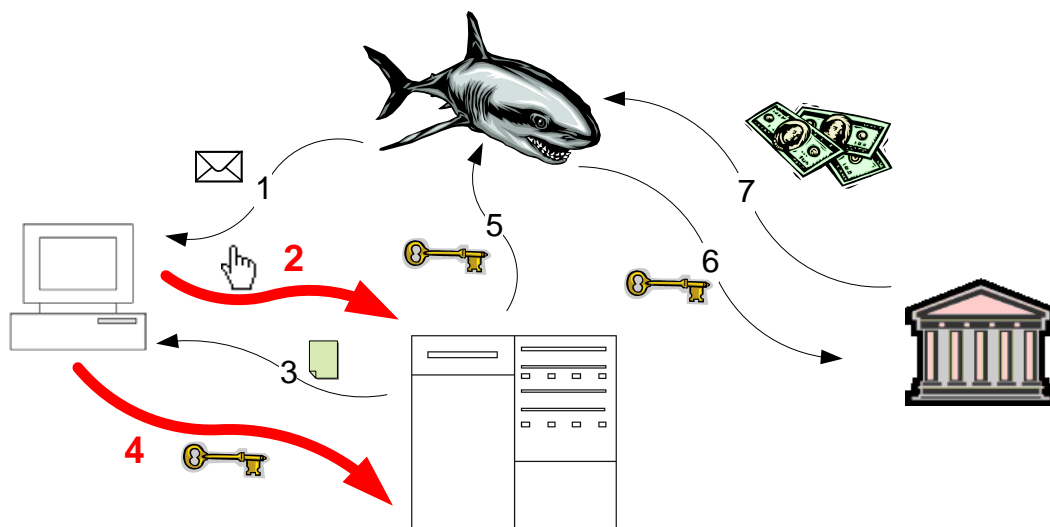
DNS-based phishing attacks rely on being able to give incorrect DNS information for a host. Since such phishing attacks rely on a user going to a site with which he or she has a previous relationship, it may be possible to detect bad information. A record could be kept, independent of the DNS cache, of previous lookups. If a name resolution yields a different result, an authoritative answer is sought from an external source known to be reliable.

It may also prove effective against attacks using numeric IP addresses (a common scenario for phishing servers on compromised home machines) to detect an access to an IP address for which no corresponding DNS lookup has been performed.

Step 2 Countermeasure: Modify Referenced Images

Phishers sometimes access images on a site controlled by the targeted company to simulate the look and feel of a legitimate email or web site. The targeted institution can detect this activity by examining the referrer field of an incoming request for an image, and once a phishing attack is underway, the web server can refuse to serve the images, or substitute the images with images displaying an informational message about the phishing attack.

This countermeasure applies to step 2, in which the image is referenced by email. It also applies to step 4, in which a web page transmitted in step 3 references imagery on a legitimate site. It can be easily circumvented by phishers hosting their own images, but has been effective in many attacks to date.



Phishing Information Flow, Steps 2 and 4

Steps 2 and 4: Prevent Navigation and Data Compromise

Step 2 in the phishing information flow is a user action that leaves the user vulnerable to a phishing attack, such as navigating to a phishing site. Step 4 is where confidential information is compromised.

Step 2 and 4 Countermeasure: Increase Inter-Application Data Sharing

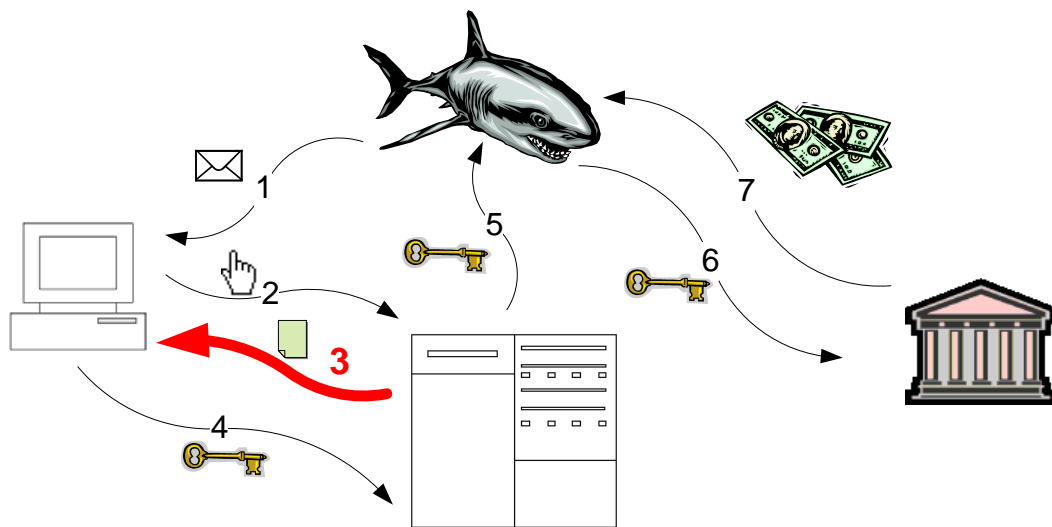
An area of future work is fighting phishing by increasing information sharing between spam filters, email clients and browsers. Important information is often lost in boundaries between these applications. A spam filter may have classified a message as likely to be illegitimate, but as long it scored below the rejection threshold, it is typically rendered by the email client on an equal basis as signed email from a trusted company.

Information gleaned while processing messages can help thwart phishing. If an email is suspicious, it can be treated differently than an authenticated message from a sender on the user's whitelist or a member of a bonded sender program.

A suspicious message can be visually indicated, scripts can be disallowed, links can be shown with their true names, forms can be disallowed, etc. This countermeasure addresses step 2 of the phishing information flow.

Similarly, once a user clicks on a link in an email message, information about the trustworthiness of the message can help determine whether to allow a traversal. Once a link is traversed, functionality (scripting, form submissions, display of links, etc.) can be restricted for links pointed to in less trustworthy messages, which can prevent step 4 in the phishing information flow from occurring.

Interfaces between spam filters, email clients and browsers that allow trustworthiness information to be transmitted would enable many new ways to combat phishing.



Phishing Information Flow, Step 3

Step 3: Preventing Transmission of the Prompt

Step 3 of the phishing information flow is a prompt to the user that will lead to the compromise of confidential information to an unauthorized party. Step 3 countermeasures attack this prompt, preventing it from arriving or from containing any leakage of information to a malicious party.

Step 3 Countermeasure: Filter Out Cross-Site Scripting

Cross-site scripting is a content injection attack that can arrive in one of two ways. A phisher can inject malicious content into a legitimate web page in step 1 of the phishing flow by storing it on the legitimate server as part of a customer review, auction, web-based email, or similar content. A phisher can also include malicious code in a URL that is included in an email to a user in step 1, for example by embedding a script in a search query that will be displayed with the search results. Such URL-embedded content will be passed from the user to the legitimate server in step 2, and returned as part of the prompt for confidential information in step 3.

Once injected, a cross-site script can modify elements of the host site so that a user believes he or she is communicating with the targeted institution, but actually is providing confidential information to a phisher.

To disrupt step 3 in the phishing information flow by preventing cross-site scripting, any user data that is ever displayed on the screen should be filtered to remove any scripts. Malicious parties have mounted cross-site scripting attacks in unexpected areas, such as date fields of web-based email pages. Rather than filtering out forbidden script elements with a “keep-out” filter, user-supplied data should be parsed with a “let-in” filter, and only permitted data elements should be allowed through.

Such filtering is a component of good web site design for independent reasons, as a cross-site script or other HTML elements could deface or alter the visual appearance of a web site, or cause other damage unrelated to identity theft.

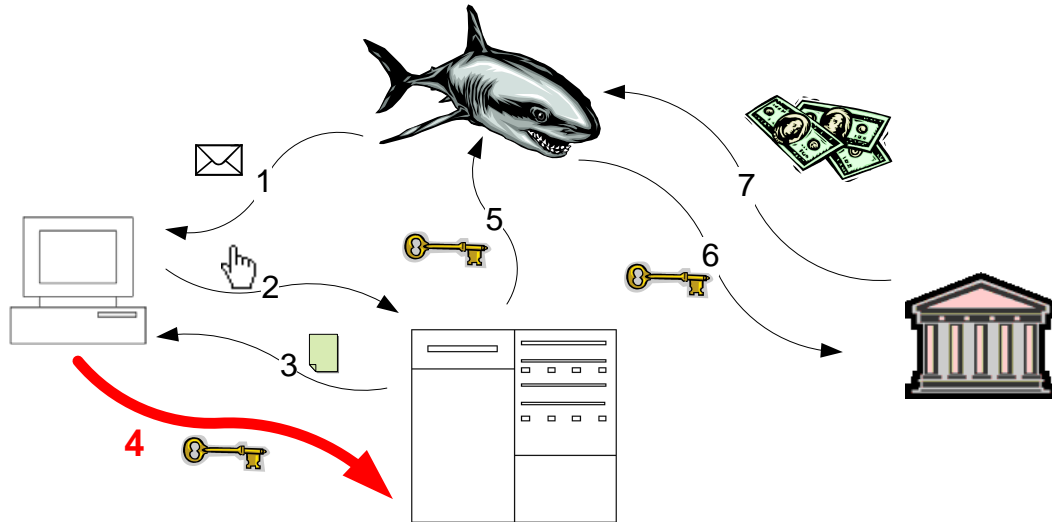
Step 3 Countermeasure: Disable Injected Scripts

There are many ways in which cross-site scripting can be introduced. It is difficult, expensive and error-prone to write an adequate filter, and often content that should be filtered is inadvertently overlooked.

A browser extension could provide protection against cross-site scripting in the future. If a new tag was introduced that could be included in HTML, such as `<noscript>`, regions could be defined in which no scripting whatsoever could occur, or in which particular functionality was prohibited. The browser could guarantee this behavior, and employing sufficient filtering would be as simple as enclosing areas of user-supplied text, such as search results or auction listings, with appropriate `<noscript>` and `</noscript>` tags.

To prevent a malicious party from including a valid `</noscript>` tag and inserting a cross-site script, a dynamically generated random key should be used that must match in the `<noscript>` and `</noscript>` tags. Such a key could be automatically generated by web content authoring tools. Since the user-supplied content would have no way to know what random number was used for the key, it would lack the information required to re-enable scripting privileges. For example:

```
[Site-supplied HTML and scripts]
<noscript key="432097u5iowhe">
[User-supplied HTML in which scripts/features are disabled]
</noscript key="432097u5iowhe">
[Site-supplied HTML and scripts]
```



Phishing Information Flow, Step 4

Step 4: Preventing Transmission of Confidential Information

Another point at which phishing attacks may be disrupted is when a user attempts to transmit confidential information at step 4 of the phishing information flow. If a deceptive phishing site can be revealed as fraudulent to the intended victim, or if the information flow can be disrupted or altered to render the confidential information unavailable or useless to the phisher, the attack can be thwarted.

In a classic deception-based phishing attack, phishers use many different techniques to maintain the deception that the user is at a legitimate site. This again involves many rapidly changing technologies. One way to deceive the user as to the location of the browser is to use deceptive links. Another is to ensure that deceptive information appears in the URL bar. For example, phishers have created Javascript programs that pop up a borderless window to obscure the real contents of the URL bar, and move the deceptive window when the user moves his browser window. Some of these Javascript programs simulate the window history if the user clicks on the history box.

It is not possible to determine whether a connection to a site is secure (i.e. uses SSL) by looking at a lock icon in a browser. There are several reasons why a lock icon cannot be trusted:


- A lock icon by itself means only that the site has a certificate; it does not confirm that the certificate matches the URL being (deceptively) displayed. A user must click on a lock icon to determine what it means, and few users ever do.
- It is possible to get some browsers to display a lock icon using a self-signed certificate (i.e. a certificate that has not been issued by a valid certificate authority), with certain encryption settings.

- A lock icon can be overlaid on top of the browser using the same technologies used to falsify the URL bar. This technique can even be used to present authentic-looking certificate data if the user clicks on the lock icon to confirm legitimacy.

While browser technologies are constantly being updated to address recent phishing tactics, browsers are large, complex programs that must provide considerable functionality and flexibility to satisfy the needs of legitimate web site designers. It is highly improbable that deceptive phishing appearances can be completely stopped solely by addressing phishing technologies piecemeal.

Step 4 Countermeasure: Anti-Phishing Toolbars

Browser toolbars are available that attempt to identify phishing sites and warn the user. These are available both as research projects and from technology suppliers. Anti-phishing toolbars use a variety of technologies to determine that they are on an unsafe site, including a database of known phishing sites, analysis of the URLs on a site, analysis of the imagery on a site, analysis of text on a site, and various heuristics to detect a phishing site. They typically display a visual indication such as a traffic light indicating the safety of a site, in which green indicates a known good site, yellow indicates an unknown site, and red indicates a suspicious or known bad site. For example:



Phishing Toolbar: eBay Account Guard

In this example, the user is viewing a page on eBay's site, so the indicator is green. Another toolbar example shows a user visiting a deceptively named site, both visually indicating the danger and providing easy navigation to a site the user most likely believes he or she is visiting:



Phishing Toolbar: Stanford SpooGuard

Anti-phishing toolbars could potentially be spoofed using current technologies. If combined with reserved screen real estate that cannot be overwritten by any page or script, this danger could be avoided.

Research has shown that users respond differently to various types of toolbar indications. In particular, a toolbar that provides specific guidance on taking or not taking an action can be over twice as effective, after user training, as a toolbar that provides only neutral or positive information about a site. Even the most effective toolbars, however, can still have phishing success rates over 10% once a user has visited a phishing site, even after the user has been trained.

Some anti-phishing toolbars use personalized information, by displaying a user-selected name or image when a user is really on a web site with which the user has a relationship.

Some browser plug-ins aim to prevent spoofing by providing a distinctive user experience for each browser window, such as animated borders or graphical patterns in a window background or in the “chrome” surrounding a browser window. The distinctive user experience is spoof-resistant because it is generated by the client on a per-session basis. Such approaches rely on the user to detect an anomalous window, and must balance the ease of detecting a spoofed window against aesthetic acceptability and intrusiveness.

Many anti-phishing toolbars go beyond presenting information about a site, and try to detect when a user is entering confidential information on a potential phishing site. The toolbar stores hashes of confidential information, and monitors outgoing information to detect confidential information being transmitted. If confidential information is detected, the destination of the information can be checked to ensure that it is not going to an unauthorized location.

Monitoring outgoing data has a challenging obstacle to overcome. Phishers can scramble outgoing information before transmitting it, so keystrokes must be intercepted at a very low level. (Some phishing toolbars wait until form submission to detect confidential information, which is ineffective against simple workarounds.) Similarly, scripts on a web page can transmit data character by character as it is typed. Moreover, some users enter keystrokes out-of-order for account and password information to avoid compromise by keyloggers, rendering even a protective keylogger ineffective. The long-term viability of outgoing data monitoring as an anti-phishing technology is unclear, but it is presently effective since most phishing attacks do not include workarounds.

Step 4 Countermeasure: Data Destination Blacklisting

Some proposals have been fielded to block data transmissions to specific IP addresses known to be associated with phishers. This is an attempt to disrupt step 4 of the phishing information flow.

Data destination blacklisting faces two major challenges. First, phishing attacks are increasingly being run in a distributed manner, using many servers in a botnet or similar configuration. It is a challenge to identify all of the phishing servers. Even if it was possible to do so, this would not prevent information

transmission in a lasting manner, as information could be transmitted through covert communications channels using the internet Domain Name System (DNS) that is used to translate host names into IP addresses. A simple example of this in which a phisher controls the DNS server for phisher.com and wants to transmit “credit-card-info” is to incur a DNS lookup on “credit-card-info.phisher.com.” The result of the DNS lookup is not important; the data has already been transmitted through the DNS request itself. Blocking DNS lookups for unknown addresses is not feasible, as DNS is a fundamental building block of the internet.

Even a blacklist that somehow managed to prevent DNS lookups on all phishing domains would still be susceptible to circumvention via DNS. Information can be transmitted via DNS even if the phisher does not control any DNS server whatsoever, by using the time-to-live fields in DNS responses from innocent third-party DNS servers.

In practice, shutting down covert communications channels is a hard problem, and it is unlikely to be effective against a determined adversary.

Step 4 Countermeasure: Screen-Based Data Entry

Some companies have deployed alternate data entry mechanisms for sensitive information. Rather than typing in the information, users enter it by selecting information on a screen. This is an attempt to disrupt step 4 in the phishing information flow for keylogging malware.

Screen-based data entry is presently effective, since phishers have not deployed workarounds. However, if screen-based data entry becomes more widely deployed, malware could intercept the display and evaluate the data displayed on the screen and the user’s interactions with it, thereby compromising the confidential information.

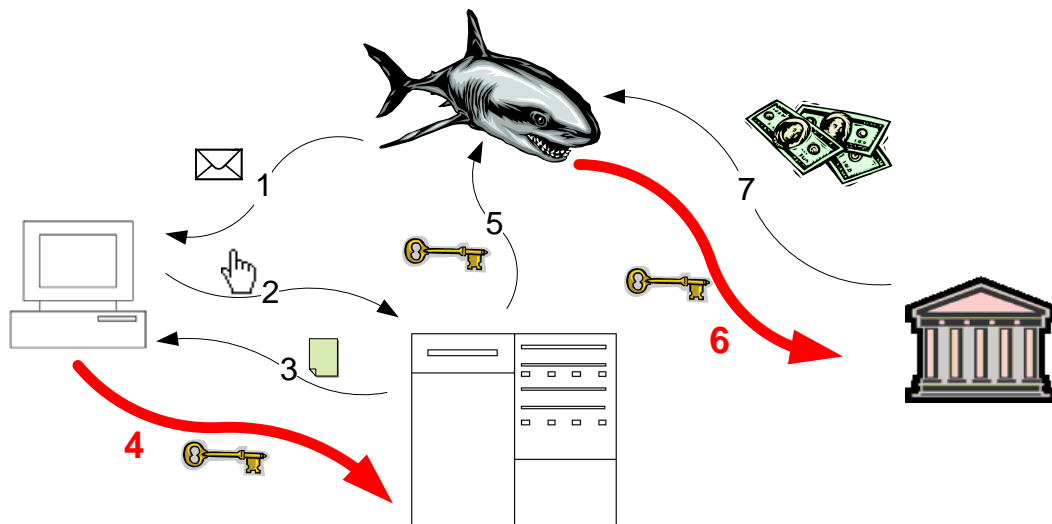
Step 4 Countermeasure: Mutual Authentication

For an authentication credential such as a password, in many cases the credential may be known to both parties. Rather than transmitting it, a mutual authentication protocol may be used to provide mutual proofs that each party has the credential, without either party ever transmitting it.

There are many such protocols. They can prove to the site that the user has the credential, while proving to the user that the site has the credential. The applicability to phishing is limited by the need to ensure that such a protocol is used – a phisher would simply ask for the credential and not run the protocol. Therefore, either all such credentials should be entered into a special program instead of a web site, or a trusted path mechanism, as discussed later, should be used. Another way to demonstrate to a user that a mutual authentication protocol will be used is to display a particular image in any window whose contents will be used for mutual authentication. Such an image is stored client-side and kept secret from external parties, so it cannot easily be spoofed.

Another potential issue with a mutual authentication protocol is that both sides must have a matching credential. It is poor practice to store passwords for users; they are normally always stored hashed with a salt. To avoid a

requirement to store interpretable passwords, a mutual authentication protocol for passwords can be combined with password hashing, which is discussed in step 6 of the phishing information flow.



Phishing Information Flow, steps 4 and 6

Steps 4 and 6: Preventing Data Entry and Rendering it Useless

Step 4 in the information flow of a phishing attack is where data is compromised, while step 6 is the use of compromised information for financial gain. Countermeasures that attack steps 4 and 6 make it less likely that information is compromised, and make it impossible for the phisher to use information in the event it is compromised.

Steps 4 and 6 Countermeasure: Trusted Path

A fundamental failing of the internet trust model is that it is not evident to a user where data being entered will ultimately be sent. A non-spoofable trusted path can ensure that sensitive information can reach only a legitimate recipient. A trusted path can protect against deception-based phishing and DNS-based phishing. If implemented in the operating system, it can also protect against application-level malware attacks.

Trusted paths have been used for login information using one of two mechanisms: a reserved area of a display, or a non-interceptable input. An example of the latter is the use of CTRL-ALT-DEL to login into a computer using an operating system in the Windows NT family, which was implemented as part of the National Computer Security Center's requirements for C2 certification.

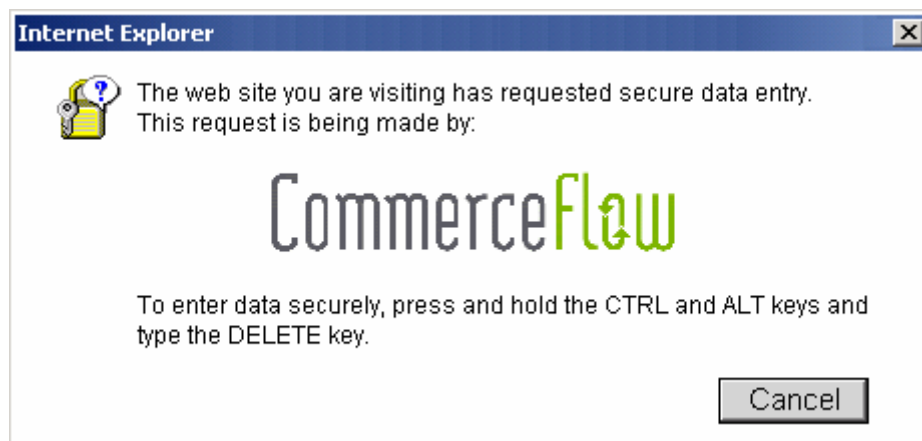
Conventional trusted path mechanisms can establish a trustworthy channel between a user and an operating system on a local machine. To be effective in combating phishing, a trusted path must be established between a user and a remote computer, over an untrusted internet, in the presence of malicious servers and proxies.

An operating system could safeguard the entry of sensitive information by providing a trusted path system service that is called with two separate types of arguments:

- A certificate, cryptographically signed by a certificate authority, which contains the identity of the requestor, a logo to be displayed and a public key; and
- Specifications for the data that is being requested.

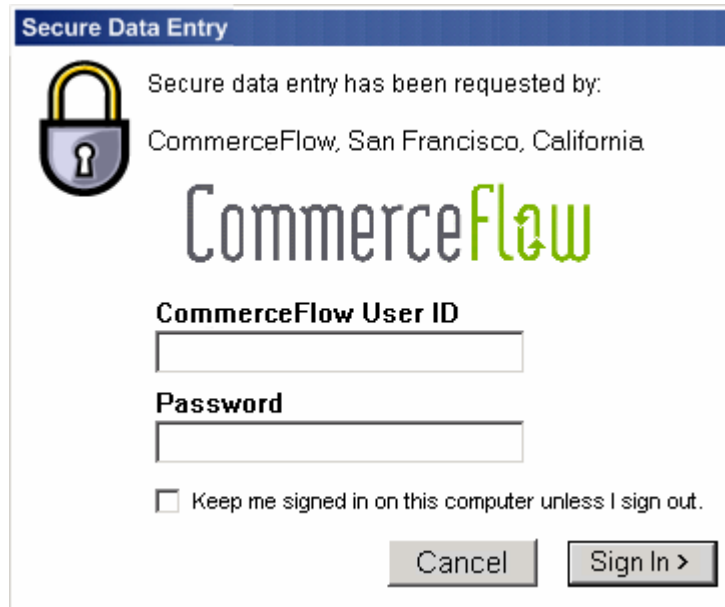
The simplest implementation of this is to use the server's certificate being used for an active SSL connection through which the current web page was received, and to have a tag in an HTML form that indicates that trusted path should be used for data input. The HTML form can be used as the specifications for requested data in a call that the browser makes to the trusted path service.

When the operating system has been notified of the impending trusted path data entry, the user is prompted to enter a non-interceptable key sequence known as a "secure attention sequence." In Windows, CTRL-ALT-DEL is a secure attention sequence. This could be used, or a potentially more user-friendly implementation is to have a special key on a keyboard dedicated to trusted path data entry.



Trusted Path: Request Notification

When the user enters the secure attention sequence, the operating system determines that trusted path data entry was requested, and displays a standard input screen, displaying the identity and logo of the data requestor from the certificate, and the specified input fields.



Trusted Path: Input Screen

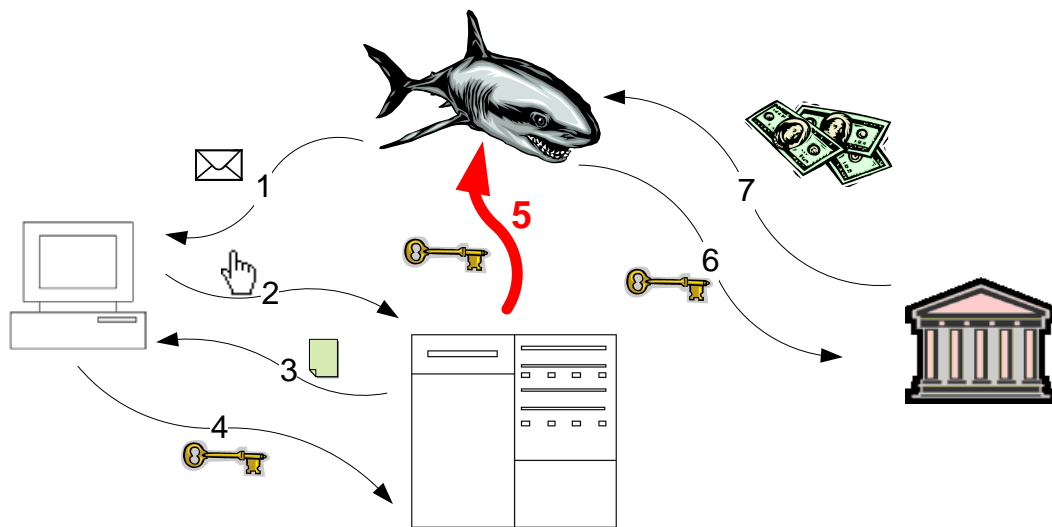
Since only the operating system will receive the secure attention sequence, the operating system is guaranteed to be in control. The trusted path data entry screen is displayed directly by the operating system in a controlled environment. In this mode, no user processes can alter the display or intercept keystrokes. This level of control by the operating system renders tampering by phishers impossible, in the absence of an administrative-level security exploit. When the fields are input, the data is encrypted by the operating system using the public key in the certificate, so that only the certified data recipient that possesses the corresponding private key can read the data. This encrypted data is then made available to the requesting application.

This particular trusted path mechanism relies on certificate authorities to verify the identity and logo of an applicant before granting a certificate. Trusted path certificates would be issued by a small, controlled set of authorities, who will require definite proof of identity and ensure that an unauthorized logo is not being used. Requirements to be a trusted certificate authority for trusted path should be at least as stringent as the requirements for root certification authorities for SSL certificates, and possibly more stringent; or a higher-security level of certificates could be required for trusted path than for SSL.

Unlike admonitions to check an advisory display element such as the lock icon, getting to a data entry screen via a trusted path is an active part of the user's interaction with a site. As users grow accustomed to always entering confidential information (passwords, credit card numbers, social security numbers, etc.) using a trusted path mechanism, any request for confidential information that does not use a trusted path would raise an immediate red flag – which could be augmented by a detection system indicating data transmission to an untrusted site, or entry of confidential information.

Trusted path is a step 4 countermeasure in that a phisher, to be able to interpret sensitive information, would need to ask for it under its actual identity, or without using a trusted path mechanism. If users are accustomed to entering sensitive information using trusted path, they are unlikely to provide it. Trusted path is also a step 6 countermeasure. A phisher could steal a certificate and ask for data using the stolen certificate. However, the phisher will be unable to interpret the sensitive data, as only the legitimate certificate owner has the private key needed to decrypt it.

Trusted path can also be implemented at an application level. The use of “@@” as a secure attention sequence for password entry in Stanford University’s PwdHash program is an application-level trusted path implementation. Trusted path implemented in the browser has the potential to protect against deception-based phishing attacks and DNS-based phishing attacks. To protect against user-privileged malware, an operating system level implementation is needed.



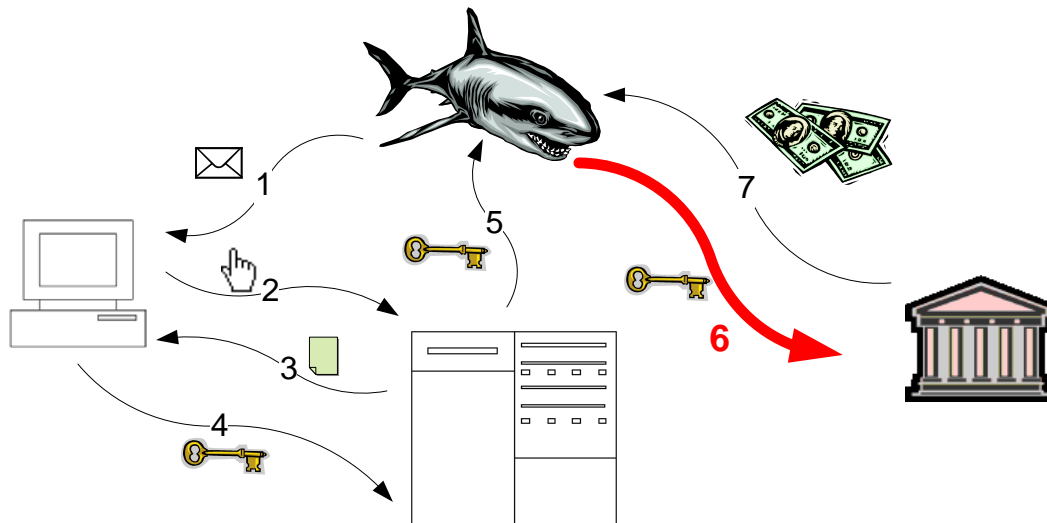
Phishing Information Flow, Step 5

Step 5: Tracing Transmission of Compromised Credentials

In step 5 in the phishing information flow, compromised credentials are obtained by the phisher from the phishing server or other collector. In the case of locally run attacks such as a Web Trojan, keylogger or local session hijacking, the phishing “server” from which the compromised credentials are obtained may be the customer’s computer.

Phishers sometimes construct elaborate information flows to cover their tracks and conceal the ultimate destination of compromised information. In some cases, these information flows span multiple media, such as compromised “zombie” machines, instant messaging, chat channels, and anonymous peer-to-peer data transfer mechanisms. Academic literature has also suggested the use of techniques such as public-key steganography, in which information could be inserted into public communications such as Usenet postings, which would make it very difficult to detect the ultimate consumer of the credentials.

In general, preventing covert communications channels is very difficult, and countermeasures at this stage generally center on taking down phishing servers before they transmit their data back to the phisher, or tracing the flow of information for prosecution of the criminals.



Phishing Information Flow, Step 6

Step 6: Interfering with the Use of Compromised Information

Another technology-based approach to combating phishing is to render compromised information less valuable. This interferes with step 6 of the phishing information flow, in which the phisher converts compromised information into illicit revenue. The following countermeasures attack step 6 of the phishing information flow.

Step 6 Countermeasure: Conventional Two-Factor Authentication

The most prevalent approach to reducing the impact of data compromise is known as *two-factor authentication*. This refers to requiring proof of two out of the following three criteria to permit a transaction to occur:

- What you *are* (e.g. biometric data such as fingerprints, retinal scans, etc.);
- What you *have* (e.g. a smartcard or dongle); and
- What you *know* (e.g. an account name and password).

Today's phishing attacks typically compromise what a user *knows*. Because such information can easily be compromised in a phishing attack, step 6 of the phishing information flow can be disrupted by requiring something a user *has* or something a user *is* in addition to password type credentials. An additional factor of authentication is commonly referred to as "second-factor authentication." Second-factor authentication can be required either to gain access to an account, or to perform a transaction. Sometimes, second-factor authentication may be required for any transaction; sometimes it may be required only for transactions

that are considered likely to be fraudulent, as discussed below under *Transaction Confirmation*.

The most widely deployed second-factor authentication device in the United States is the one-time-passcode (OTP) device. Such a device displays a code that changes either on regular intervals, or each time it is used. To demonstrate that a user has the device, the user is prompted to enter the current passcode, which is validated by a server that knows the sequence that is used and the current value.

OTPs are easy to understand, and they can largely remove the secondary market for the subsequent sale of stolen credentials. However, they are vulnerable to phishing attacks in which the economic damage takes place while the OTP is still valid. If a short period of time elapses between the time the OTP is given to the phisher in step 4 of the phishing information flow and the use of the credential in step 6, which can for example be the case in a man-in-the-middle attack or a session hijacking attack, then the phisher can use the OTP in step 6.

Other forms of second-factor authentication are resistant to such attacks. Smart cards and USB dongles can perform onboard cryptographic processing and ensure that they are authenticating directly to an authorized party, in a manner that an eavesdropper would be unable to interpret. Well-implemented biometric authentication systems use a challenge-response protocol that is bound to the communications channel in such a way that a man-in-the-middle cannot reuse responses to challenges posed by the ultimate server.

Step 6 Countermeasure: Computer-Based Second-Factor Authentication

Separate hardware second-factor authentication devices can be an effective countermeasure. However, they are expensive to purchase, deploy and support, and some (such as smart cards) require daunting infrastructure investments. Additionally, customers have been resistant to using hardware second-factor authentication devices, due to the inconvenience that can be involved. Conventional two-factor authentication is appropriate for high-value targets such as commercial banking accounts, but so far has not been widely deployed in the United States for typical consumer applications.

A less costly approach to second-factor authentication is to use the customer's computer as a *what you have* authentication factor. This is based on the observation that customers typically perform their online banking from one of a small number of home or work computers. Computer-based second-factor authentication registers those authorized computers with the customer's account, and uses their presence as a second-factor authentication.

This is a valid approach, and has significant cost and usability advantages when compared with hardware-based second-factor authentication. However, there are some security considerations. First, the identity information for the machine should be transmitted in a manner that is not susceptible to a man-in-the-middle attack, such as using a special software program that authenticates the recipient

of the identity information, or a secure cookie that will be sent only to a remote site that has authenticated itself using SSL, to avoid DNS-based attacks receiving the authentication information.

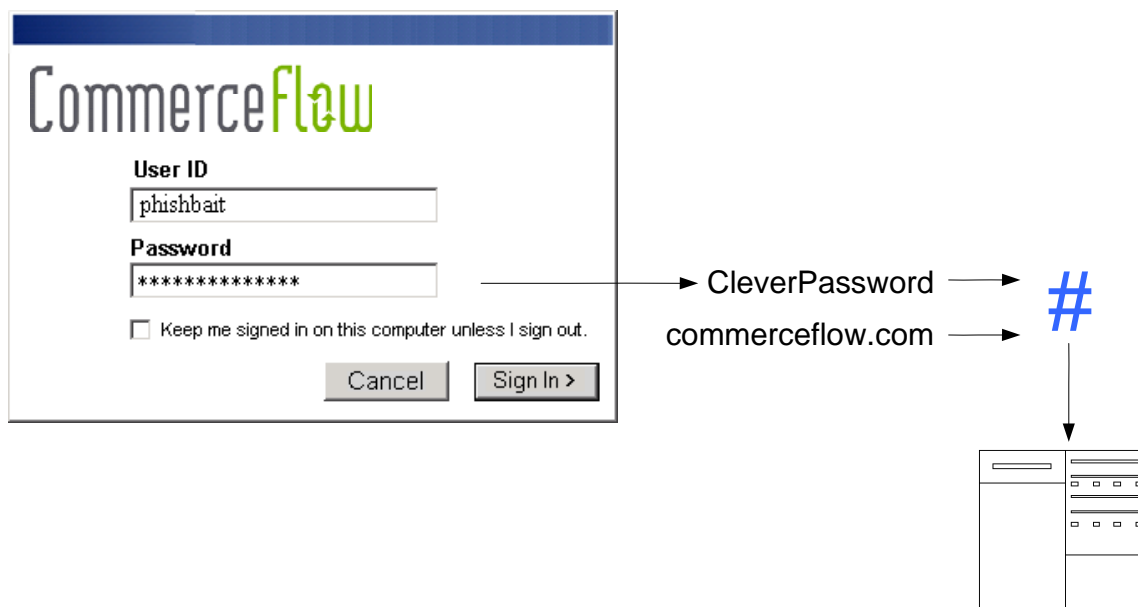
Second, computer-based authentication may ultimately be susceptible to a locally running session hijacking attack or other attack in which transactions are performed using the user's computer. In some sense, once malicious software is executing on a customer's computer, that computer is no longer something the customer has, and is rather something the phisher has, and can use for authentication.

A key security issue in computer-based second-factor authentication is the authorization of new computers, or re-authorization of existing computers. Users may sometimes need to use a foreign or newly obtained computer, or to reauthorize a computer when its authorization information has been removed. Computer authorization is sometimes performed by asking the user to answer some questions or provide a secondary password. This information can potentially be phished, and as a second *what you know* factor, reduces computer-based authentication to single-factor authentication.

To be a true second authentication factor, computer-based authentication needs to use a *what you have* to authorize a new computer. For example, when a user requests an authentication of a new computer, a one-time passcode can be sent to his cell phone. The user can then type this information into a special program that will send it only to the proper destination. It is important in this case that the user never enters the passcode into a web page, as a phishing site could obtain the passcode and use it to authorize a phishing machine. Another form of *what you have* for computer authorization is a clickable authorization link in an email, which authorizes a computer at the IP address used to click on the link.

Step 6 Countermeasure: Password Hashing

Phishing for passwords is worthwhile only if the password sent to the phishing server is also useful at a legitimate site. One way to prevent phishers from collecting useful passwords is to encode user passwords according to where they are used, and transmit only an encoded password to a web site. Thus, a user could type in the same password for multiple sites, but each site – including a phishing site – would receive a differently encoded version of the password. An implementation of this idea is called password hashing. In password hashing, password information is hashed together with the domain name to which it is going before it is transmitted, so that the actual transmitted passwords can be used only at the domain receiving the password data. Password hashing could ultimately be provided by a browser as a built-in mechanism that is automatically performed for password fields. To prevent offline dictionary attacks, a site using password hashing should also enforce good password requirements. Password hashing is a countermeasure to step 6 of the phishing information flow, in that password data compromised in a deception-based phishing attack cannot be reused on the legitimate site.



Password Hashing

In addition to phishing security, password hashing provides good protection against non-phishing forms of identity theft based on large-scale theft of password data from a site. It provides assurance both that sites will not store plaintext password data, and that the passwords cannot be reused on another site. Users commonly use the same password on multiple sites, so stolen user name and password data from one site can be reused on another. As long as passwords are difficult to guess through a dictionary attack, password hashing prevents such cross-site reuse of stolen credentials. Password hashing can also be combined with a mutual authentication protocol to obviate the need to store a mutually authenticated password in plaintext.

By itself, password hashing does not provide protection from a deceptive phishing attack, as a phisher would not perform the password hashing after asking for a password. Therefore, a way to make password entry different from other data entry to enforce password hashing is needed. A trusted path, as discussed earlier, is appropriate for this purpose. Stanford University's PwdHash program uses "@@" as a secure attention sequence to ensure that password hashing is used for an input field. This secure attention sequence is intercepted by a browser plug-in that obscures the password data to keep it from scripts until focus leaves the field or a form is submitted, at which point a hashed version of the password is substituted.

Step 6 Countermeasure: Transaction Confirmation

One approach to reducing phishing risk is to concentrate on online transactions that may be fraudulent. This is analogous to the risk management measures that banks take in the physical world: every credit card transaction is evaluated, and suspicious transactions are checked with the customer.

An analysis of online transactions may be performed using a variety of metrics, such as the user's IP address, the presence of authentication information such as a cookie on the user's machine, the amount of a transaction, the destination bank account, characteristics of the destination bank account, and cross-account analysis of transactional patterns. Such analysis can be performed by software integrated into a bank's online systems, or by an "appliance" that monitors web traffic. When a transaction is flagged as suspicious, transaction-specific authentication from the customer is required.

Critically, such authentication should not be in the form of "what you know" questions that can be phished. A strong form of transaction authentication uses a trusted device as a second factor, such as a telephone. If a phone call is placed to a customer at a number that is known to belong to him or her, or an SMS message is sent to the customer's cell phone, the customer can confirm the transaction by voice or return message. It is important that confirmation information includes details of the transaction itself, since otherwise a phisher could perform a session hijacking attack and alter a transaction that the user will confirm. Biometrics could also be used for authentication, provided that the biometric device had a way to trustably display transaction details.

Some research indicates that customers may confirm transactions without checking the details, if they are expecting to have to confirm. Therefore, such confirmations should be very unusual, or a user interface should be used that requires that the user actively select a transaction to confirm.

Transaction analysis and confirmation, when implemented well, is an effective step 6 mitigation across all types of phishing fraud, including administrative-privileged malware, as well as other forms of non-phishing identity theft. It does not provide 100% protection, but can significantly reduce losses through online transactions. Banks should evaluate this benefit against deployment costs and potential user experience disruptions.

Step 6 Countermeasure: Policy-Based Data

Another step 6 countermeasure is to render data unusable to a third party by inextricably combining it with a policy that dictates how or by whom the data can be used. This is not only a countermeasure against step 6 of the phishing attack, and also can be applied to non-phishing identity theft such as data theft by hacking or insider compromises.

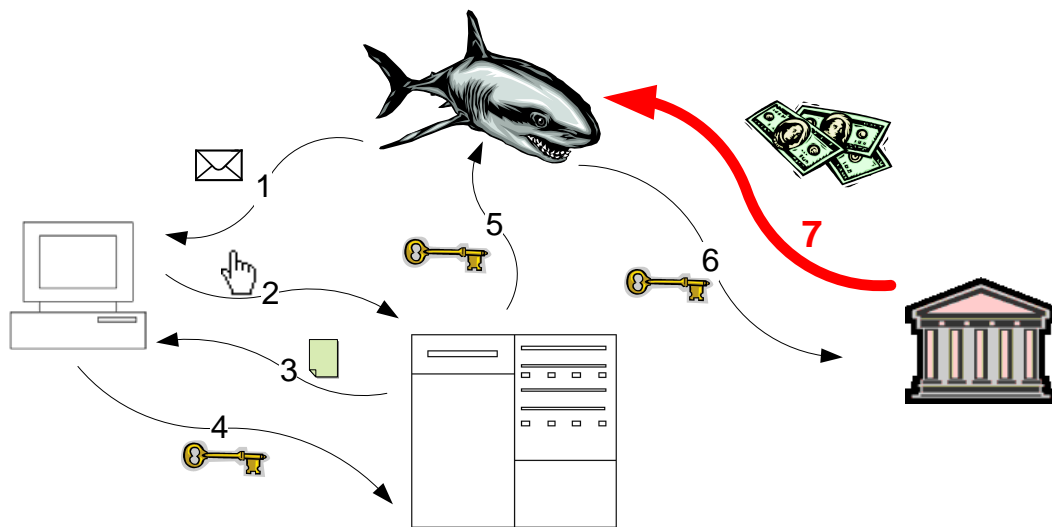
This technique is appropriate in situations in which the site receiving and storing the data is not the ultimate consumer of the data. For example, e-commerce sites and ISPs need to keep credit card numbers on file so users can conveniently charge their purchases or pay recurring bills. However, the credit card transaction is not performed by the e-commerce company or ISP. A payment processor is actually responsible for making the charge.

A site can combine the credit card information with a policy that stipulates that only that site can make a charge. This combined information is then encrypted using a public key belonging to the payment processor, before the credit card

information is stored. The information cannot be decrypted without the private key, which only the payment processor has. So even if the data is stolen, it is useless to the thief. This differs from conventional encrypted database schemes in that normally, someone at the company storing the data has access to the decryption key, and such people can be bribed, or the decrypted data can otherwise be compromised.

When a transaction is to be performed, the encrypted credit card information is sent along with the transaction details. The payment processor decrypts the bundle and checks the policy. If the transaction is not authorized under the policy – for example if the policy says that only CommerceFlow is allowed to charge the card, while PhishingEnterprises is attempting a charge – then the charge is declined.

To be a phishing countermeasure, this can be combined with the trusted path mechanism discussed earlier. A policy can be embedded in a form, combined with specified input fields and encrypted with the specified key, which can have identity information displayed on-screen. Even if a phisher was somehow able to gain access to the site's key data or otherwise compromise the information, the policy-based credentials would remain useless to the phisher.



Phishing Information Flow, Step 7

Step 7: Interfering with the Financial Benefit

In step 7 of the phishing information flow, economic gain is realized from the use of the compromised credentials in step 6.

Financial institutions have instituted delays in certain types of money transfers to allow for detection of accounts used in a phishing attack. If a fraudulent receiving account is identified during the holding period, the transfer can be voided and the economic damage can be averted.

Step 7 is also of considerable interest to law enforcement. Often, phishers can be caught by tracing the flow of money from the use of stolen credentials.

Phishers often filter money through multiple layers and use anonymous cash instruments, but ultimately a phisher receives the money, and such a flow can often be traced.

Non-Technical Best Practices

This report is primarily concerned with anti-phishing technologies. Nonetheless, there are some practices that any potential phishing target should be aware of:

- Register the most deceptive available domain names similar to your brands. This is the cheapest insurance you can buy.
- Trademark your domain names to provide recourse against a party who registers deceptively similar domain names.
- Monitor recent domain registrations and take action against parties registering domain names deceptively similar to yours.
- Publish email authentication information in DNS records, and use authenticated email for all customer communications. This should include any parties who send mail on your behalf.
- Consider digitally signing all outgoing emails to your customers. This can be performed at an email gateway if it is not feasible to do so on your mail servers.
- Establish clear policies on your email practices, such as never asking for personal information or possibly never providing a clickable link in an email. Be sure that your policies are acceptable to all stakeholders in your organization. Enforce your policies with all third parties that send email on your behalf. Communicate your policies to your customers regularly, preferably in every email communication and in other media, such as printed statements.
- Include personalized information in each email to a customer. Along with the personalized information, include an educational statement that it is your policy always to do so.
- Provide an email address such as spoofof@yourcompany.com, which customers may submit an email to and determine whether the email is legitimately from you or not. Provide clear instructions on your web site, and in communications from your company, on how to report a phishing message.
- Do not use web sites with unusual or unpredictable names for customer interactions.
- Ensure that your web site uses SSL and that all certificates are current.
- Remove any open URL redirects from your site.
- Ensure that all user-supplied data is stringently filtered, using a let-in filter, for cross-site scripting and SQL injection.

- Institute a senior position in your organization with responsibility for identity theft losses, whose responsibilities do not also include other potential losses (such as bad credit decisions) that could distract attention from phishing losses.
- Establish a cross-functional task force responsible for responding to phishing attacks. Personnel involved should be senior and empowered to make and implement decisions quickly. Clearly delineate responsibilities and procedures. Hold “fire drills” to ensure that roles are understood and hand-offs are smooth.
- Proactively prepare customer communications to be sent out in the event of a phishing attack, to avoid delays in sending them when an attack is underway.
- Monitor signs of a phishing attack, including email bounce messages, customer call volumes, anomalous account activity, suspicious image use of images, discussions on phishing groups, etc.
- Notify email filtering companies that use signature-based checking immediately when a phishing attack is underway and provide them with samples of the phishing emails. Such companies may be able to deploy rules that will block many emails from reaching their intended recipients.
- Notify law enforcement promptly when a phishing attack is confirmed. (See Appendix B.)
- When a phishing attack is confirmed, post an alert on your web site and consider informing customers of the attack via email.
- Trace the phishing servers and get them shut down as quickly as possible. Service providers are available that can assist in this effort.
- Staff up your customer service when a large-scale phishing attack is confirmed.
- Preserve evidence of the phishing attack for subsequent prosecution of the phishers.
- Do not empower any third parties to act on your behalf in violation of the preceding practices.

Conclusions

No single technology will completely stop phishing. However, a combination of good organization and practice, proper application of current technologies, and improvements in security technology has the potential to drastically reduce the prevalence of phishing and the losses suffered from it. In particular:

- High-value targets should follow best practices and keep in touch with continuing evolution of them.
- Phishing attacks can be detected rapidly through a combination of customer reportage, bounce monitoring, image use monitoring, honeypots, alert administrative action and other techniques.
- Email authentication technologies such as Sender-ID and cryptographic signing, when widely deployed, have the potential to prevent phishing emails from reaching users.
- Analysis of imagery is a promising area of future research to identify phishing emails.
- Encrypted patches could prevent malware authors from knowing about security vulnerabilities, and automate rapid responses to threats.
- Personalized information should be included in all email communications. Systems allowing the user to enter or select customized text and/or imagery are particularly promising.
- Browser security upgrades, such as displaying potentially deceptive content distinctively and warning when a potentially unsafe link is selected, could substantially reduce the efficacy of phishing attacks.
- Detecting fraudulent DNS information is a promising area of investigation.
- Information sharing between the components involved in a phishing attack – spam filters, email clients and browsers – could improve identification of phishing messages and sites, and restrict risky behavior with suspicious content.
- Content injection attacks are an increasing problem. All user content should be filtered using a let-in filter. Browser security enhancements could decrease the likelihood of cross-site scripting attacks.
- Anti-phishing toolbars are promising tools for identifying phishing sites and heightening security when a potential phishing site is detected.
- Domain-specific modification of credentials, such as password hashing, is a powerful identity theft countermeasure, and can be a highly effective anti-phishing measure when combined with a trusted path.
- An OS-level trusted path for secure data entry and transmission has the potential to dramatically reduce leakage of confidential data to unauthorized parties.

- Hardware-based two-factor authentication is highly effective against phishing, while some approaches are susceptible to short-term phishing attacks, and is recommended in situations in which a small number of users are involved with a high-value target.
- Computer-based two-factor authentication offers good security characteristics and low deployment cost, while being susceptible to certain types of malware attacks. A key security design consideration is how to authorize a new computer.
- Transaction confirmation through an out-of-band communications channel is a powerful technique that is resistant to malware.
- Whenever possible, data should be combined with a policy statement that would prevent use by a phisher, and encrypted with a public key belonging to a downstream data consumer. This can prevent unauthorized use of compromised data.
- Ensuring that logs are kept that can be of assistance with law enforcement, and being able to quantify a loss, can render a law enforcement response much more effective.

Appendix A: Technology Vendors

The vendors in this appendix are representative providers of anti-phishing technology and services. This appendix is provided for informational purposes only. The United States Department of Homeland Security cannot ensure that this list is complete or correct, and does not endorse any specific vendor.

Monitoring, Alarming, Investigation & Takedown

This category covers solutions that monitor activities on the network and raise an alarm when a potential phishing attack is either being prepared or is in progress. Takedown activities can then occur to bring the phishing site down.

Companies in this category may provide a wide range of services, using a variety of different approaches.

- OSpam.net
- Brandimensions
- Corillian
- Cyota
- Cyveillance
- ICG
- iDEFENSE
- Internet Identity
- MarkMonitor
- NameProtect
- Netcraft
- SAIC
- Secure Science
- Verisign

Helping a consumer to identify a financial institution

Encrypted Email

- Alien Camel
- PostX
- Sigaba
- Tumbleweed

Email Containing Personalized information

- PassMark Security

Email Filtering to Remove Fraudulent Email

- 0Spam
- Brightmail
- Cloudmark
- Digital Envoy
- Engate
- Ironport
- MailChannels
- MailFrontier
- MessageLevel
- Sigaba

Identifying a Valid Web Site

- Billeo
- PassMark Security
- Stanford SpoofGuard
- Whole Security

Providing stronger authentication

Two-Factor Authentication

- Software PKI Certificates
 - GeoTrust
 - Thawte
 - Verisign
- Hardware tokens
 - Aladdin
 - Authenex
 - CryptoCard
 - Kobil Systems
 - RSA SecureID
 - SafeNet iKey

- Secure Computing
- Thales
- Vasco
- Smart Cards
 - ActivCard
 - AOS-Hagenuk
 - Gemplus
 - Kobil Systems
 - Thales
 - Todo
 - Vasco
 - Xiring
- Virtual Second-Factor Authentication
 - Anakam
 - Arcot Systems
 - iBIZ
 - PassMark Security
 - TriCipher
 - WiKID Systems
- Biometric
 - Bio-Key International
 - Bioscrypt
 - DigitalPersona

Transactional Authentication

- 41st Parameter
- Cydelity
- Cyota eSphinx

Desktop Technologies

Toolbars and Phishing Site Detection

- Billeo
- Cloudmark

- Comodo
- CoreStreet
- Deepnet Technologies
- GeoTrust
- Netcraft
- Stanford SpoofGuard
- Websense
- Whole Security

Malware Detection

- F-Secure
- Internet Security Systems
- McAfee
- Panda Software
- Spybot Search & Destroy
- Sophos
- Symantec
- WebRoot
- Websense
- Whole Security

Educational and Consulting Services

- Glennbrook Partners
- Internet Identity
- Radix Labs
- Secure Science

Appendix B: Law Enforcement Resources

Consumers receiving a phishing email should report the email to the institution being targeted.

A business, if victimized by a phishing attack, is encouraged to contact a law enforcement agency – local, state or federal – to pursue an investigation or other appropriate response. There are a number of state and local high tech crimes units that are appropriate. Due to the global nature of many of these attacks, the unit should have experience investigating crimes in other countries and jurisdictions.

The United States Secret Service, through its Field Offices, Electronic Crimes Working Groups and sixteen Electronic Crimes Task Forces nationwide, has particular expertise in investigating phishing attacks. Secret Service field offices may be found at http://www.uss.s.treas.gov/field_offices.shtml.

The Federal Bureau of Investigation (FBI) has wide-ranging expertise in identity theft cases. Phishing attacks should be reported to the FBI through the Internet Fraud Complaint Center at <http://www.ifccfbi.gov/index.asp>.

A victimized business should also report phishing attacks to the Federal Trade Commission (FTC). A form for submitting a report to the FTC may be found at <http://www.consumer.gov/idtheft>.

The following actions will assist law enforcement in an investigation:

- Preserve all log data.
- Have consumers forward phishing e-mail, complete with header information, as well as any information they provided to the bogus request. This information is essential in tracing the e-mail route, ensuring the preservation of evidence and providing law enforcement with verifiable information for comparison.
- Record the level of returned or bounced e-mails to assist in estimating the scope of the attack.
- Provide as much information on the phishing IP addresses as available, and coordinate any attempts or efforts to persuade the Internet Service Provider to shut down the illegitimate website with law enforcement. In some instances, the site may need to be left up a short time to assist law enforcement in pinpointing the origin and gathering as much information as available to aid in identifying the origination location.
- Provide information on compromised customers who are willing to cooperate with a law enforcement investigation by providing account numbers, locations, etc.

Appendix C: Phishing Bibliography

Ben Adida, David Chau, Susan Hohenberger and Ronald L. Rivest, *Lightweight Signatures for Email*. Draft of June 18, 2005; to appear.

Steven M. Bellovin, *Using the Domain Name System for System Break-ins*. Proceedings of Fifth Usenix UNIX Security Symposium, June 1995.

William E. Burr, Donna F. Dodson and W. Timothy Polk, *Electronic Authentication Guideline: Recommendations of the National Institute of Standards and Technology*. NIST Technology Administration, US Department of Commerce, Special Publication 800-63, September 2004.

N. Chou, R. Ledesma, Y. Teraguchi and J.C. Mitchell, *Client-Side Defense Against Web-Based Identity Theft*, 11th Annual Network and Distributed System Security Symposium (NDSS '04), San Diego, February 2004.

Tyler Close, *Waterken YURL: Trust Management for Humans*. Waterken Technical Report, July 2004.

Fred Cohen, *50 Ways to Attack Your World Wide Web System*. Computer Security Institute Annual Conference, Washington, DC, October 1995.

F. De Paoli, A. L. DosSantos and R. A. Kemmerer. *Vulnerability of "Secure" Web Browsers*. Proceedings of the National Information Systems Security Conference, 1997.

Rachna Dhamija and J. D. Tygar. *The Battle Against Phishing: Dynamic Security Skins*. Proceedings of the 2005 ACM Symposium on Usable Security and Privacy.

Aaron Emigh, *Phishing Countermeasures*. DIMACS Workshop on Theft in E-Commerce: Content, Identity and Services, April 14, 2005.

Aaron Emigh, *Trusted Path in Heterogeneous Environments*. First Workshop on Trustworthy Interfaces for Passwords and Personal Information, June 13, 2005.

Aaron Emigh and John Mitchell, *Anti-Phishing Technology*. Report of the US Secret Service San Francisco Electronic Crimes Task Force, January 19, 2005.

Federal Deposit Insurance Corporation, *Putting an End to Account-Hijacking Identity Theft*. FDIC Division of Supervision and Consumer Protection, Technology Supervision Branch, December 14, 2004.

Edward W. Felten, Dirk Balfanz, Drew Dean, and Dan S. Wallach, *Web Spoofing: An Internet Con Game*. 20th National Information Systems Security Conference (Baltimore, Maryland), October 1997.

Financial Services Technology Consortium, *Understanding and Countering the Phishing Threat*. FSTC Counter-Phishing Project Whitepaper, January 31, 2005.

Amir Herzberg and Ahmad Gbara, *TrustBar: Protecting (even Naïve) Web Users from Spoofing and Phishing Attacks*. Draft of November 11, 2004; forthcoming.

The HoneyNet Project & Research Alliance, *Know Your Enemy: Phishing*. Report of the HoneyNet Project, May 16, 2005.

IEEE P1363 Working Group, *IEEE P1363.2: Standard for Password-Based Public Key Cryptographic Techniques*. IEEE Draft D20, March 28, 2005; forthcoming.

Markus Jakobsson, *Modeling and Preventing Phishing Attacks*. Phishing Panel in Financial Cryptography '05.

Audun Jøsang and Mary Anne Patton, *User Interface Requirements for Authentication of Communication*. Proceedings of the Fourth Australian User Interface Conference on User Interfaces, Volume 18, February 2003.

Dan Kaminsky, *Black Ops of DNS*. Black Hat Briefings 2004.

Avivah Litan, *Phishing Attack Victims Likely Targets for Identity Theft*, Gartner FirstTake FT-22-8873, May 4, 2004.

V. Benjamin Livshits and Monica S. Lam, *Finding Security Vulnerabilities in Java Programs Using Static Analysis*. 14th Usenix Security Symposium, August 2005.

Robert T. Morris, *A Weakness in the 4.2BSD UNIX TCP/IP Software*. Computing Science Technical Report 117, AT&T Bell Laboratories, February 1985.

Anton Rager, *Advanced Cross-Site-Scripting with Real-time Remote Attacker*. Avaya Labs Technical Report, February 9, 2005.

Rod Rasmussen, *Phishing Prevention: Making Yourself a Hard Target*. Internet Identity / APWG, April 5, 2004.

Eric Rescorla, *Optimal Time to Patch Revisited*. RTFM.com working paper.

Blake Ross, Collin Jackson, Nick Miyake, Dan Boneh and John C. Mitchell, *Stronger Password Authentication Using Browser Extensions*. Proceedings of the 14th Usenix Security Symposium, 2005.

RSA Security, *3rd Annual Opinion Research Corporation Security Survey*. RSA Security Report, February 14, 2005.

Joseph Stewart, *Win32.Grams E-Gold Account Siphoner Analysis*. LURHQ Threat Analysis Report, November 4, 2004.

Min Wu, Robert Miller and Simpson Garfinkel, *Do Security Toolbars Actually Present Phishing Attacks?* Symposium On Usable Privacy and Security 2005.

Min Wu, Robert Miller and Simpson Garfinkel, *Secure Web Authentication with Mobile Phones*. DIMACS Symposium On Usable Privacy and Security 2004.

Zishuang (Eileen) Ye and Sean W. Smith, *Trusted Paths for Browsers*, 11th Usenix Security Symposium, August 2002.

Zishuang (Eileen) Ye, Y. Yuan and Sean W. Smith, *Web Spoofing Revisited: SSL and Beyond*. Technical Report TR2002-417, Department of Computer Science, Dartmouth College. February 2002.

Appendix D: Other Resources

The Anti-Phishing Working Group (APWG)	http://www.antiphishing.org
Financial Services Technology Consortium	http://www.fstc.org
The Internet Fraud Complaint Center	http://www.ifccfbi.gov/index.asp
The Identity Theft Data Clearinghouse	http://www.consumer.gov/idtheft
Digital PhishNet	http://www.digitalphishnet.org

Appendix E: The DHS-SRI Identity Theft Technology Council

This document is a publication of the DHS-SRI Identity Theft Technology Council (ITTC). The ITTC is studying a range of anti-phishing defenses and exploring ways that various technologies may be deployed, improved, and combined. This report summarizes the working findings as of the document release date. The members of the ITTC are:

Val Batiste	Wells Fargo Bank
Jennifer Bolt	Franklin Templeton
Dan Boneh	Stanford University
Jeffrey Camiel	Information INC.
Ebrima Ceesay	UC Davis
Skylar Cranmer	UC Davis
Mary Ann Davidson	Oracle
Drew Dean	SRI International
Dean DeBiase	Startup Partners
Aaron Emigh	Radix Labs
Ian Foley	SRI International
Bill Harris	Passmark Security
Nicholas Imparato	Hoover Institute
David Jevans	Anti-Phishing Working Group
Chris Lalonde	eBay
Karl Levitt	UC Davis
Patrick Lincoln	SRI International
Ulf Lindqvist	SRI International
Robert Lozito	Sacramento County Sheriff's Department
Douglas Maughan	US Department of Homeland Security
John Mitchell	Stanford University
Peter G. Neumann	SRI International
Richard Perlotto	Cisco Systems
Robert D. Rodriguez	Robert D. Rodriguez and Associates, LLC
Jim Roskind	Radix Labs

Jeff Rowe	UC Davis
Eduardo Roy	Squires, Sanders & Dempsey, LLP
Marc Sachs	SRI International
Abe Smith	Xilinx
Jeff Smith	Tumbleweed
Tye Stallard	UC Davis
Wen Tseng	Washington Mutual
Doug Tygar	UC Berkeley
Chris Von Holt	US Secret Service
Don Wilborn	US Secret Service