

Anti-Phishing Technology

This report was produced in conjunction with the United States Secret Service
San Francisco Electronic Crimes Task Force

Contact: Aaron Emigh
Radix Labs
ate@radixlabs.com

January 19, 2005

Executive Summary

This report focuses on technology countermeasures to a form of electronic identity theft commonly known as “phishing.” While there is no single silver bullet that can stop all phishing attacks, a combination of existing technology and practices can substantially reduce the risk and financial impact of phishing. Future technologies hold the promise of dramatic security improvements.

Introduction

Phishing is a form of identity theft in which deception is used to trick a user into revealing confidential information with economic value. Similar forms of identity theft, in which worms or viruses install keyloggers, are sometimes also referred to as phishing. This report focuses on phishing involving deceptive electronic messages.

While the term “phishing” originated in AOL account theft using instant messaging, the most common type of phishing message today is email. In a typical scenario, a phisher sends fraudulent email, in bulk, claiming that there is a problem with a recipient’s account at a financial institution or other business. The email asks the recipient to visit a web site and provides a link. If a recipient enters a valid user name and password into the fraudulent web site, the phisher can impersonate the victim. This may allow the phisher to transfer funds from the victim’s account or cause other damage.

There are many variations on this scheme. It is possible to phish for other information in addition to user names and passwords, such as credit card numbers, bank account numbers, social security numbers or mothers’ maiden names. With HTML email readers, it is also 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. In browser-based attacks, it is possible to use Javascript to take over the address bar or otherwise deceive the user into believing he or she is communicating with a legitimate site.

Phishing presents direct risks through the use of stolen credentials and indirect risk to institutions that conduct business on line through erosion of customer confidence.

The frequency of phishing attacks has increased dramatically in recent months, as has the sophistication of attacks.

The Gartner group estimates the direct phishing-related loss to US banks and credit card issuers in the last year to be \$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 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 a cost-effective investment.

Phishing Technologies

Phishers use a wide variety of technologies, with one common thread. All technologies employed by phishers have the goal of deception. For example:

- Deceiving a user into believing a message comes from a trusted source;
- Deceiving a user into believing that a web site is a trusted institution;
- Deceiving a spam filter to classify a phishing email as legitimate.

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. As financial institutions have increased their online presence, the economic value of compromising account information has increased dramatically. Criminals such as phishers can afford to invest in technology commensurately 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. Given that, a brief review of typical practices will help illuminate the problem, and motivate the countermeasures. Descriptions of recent phishing attacks, and related statistics, may be found at <http://www.anti-phishing.org>.

Deceptive return address information

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 an official-looking domain (e.g. “commerceflow-security.com” to spoof a company whose real domain is “commerceflow.com”) and sending email from that domain name.

Fraudulent request for action

Phishing emails typically claim to require user action to prevent a problem with an account. They may claim to be conducting a security audit, or to have

detected fraud on the user's account, or to require updated contact information. The request for action must seem authentic to prompt the user to reveal confidential information.

Deceptive appearance

Effective phishing emails and web sites must present a visual appearance consistent with the institutions that they are mimicking. There are many elements to this.

Deceptive visual appearance

Color schemes and imagery mimic the targeted institution.

Deceptive links

A call to action from a phisher typically requires a user to click on a link in a message to go to a web site. The phisher's web site does not have a legitimate name, so the actual destination is often disguised. (An exception to this rule is web sites that are simply misleadingly named.) 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 their 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.

Deceptive location

Once a phisher has convinced a user to click on a link, the phisher must maintain the deception that the user is at a legitimate location. This again involves many rapidly changing technologies. One aspect of deceiving 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 window when the user moves his 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 ever do.
- It is possible to get a browser 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 may be overlaid on top of the browser using the same technologies used to fake the URL bar. This technology may even be used to present authentic-looking certificate data if the user clicks on the lock icon to confirm legitimacy.

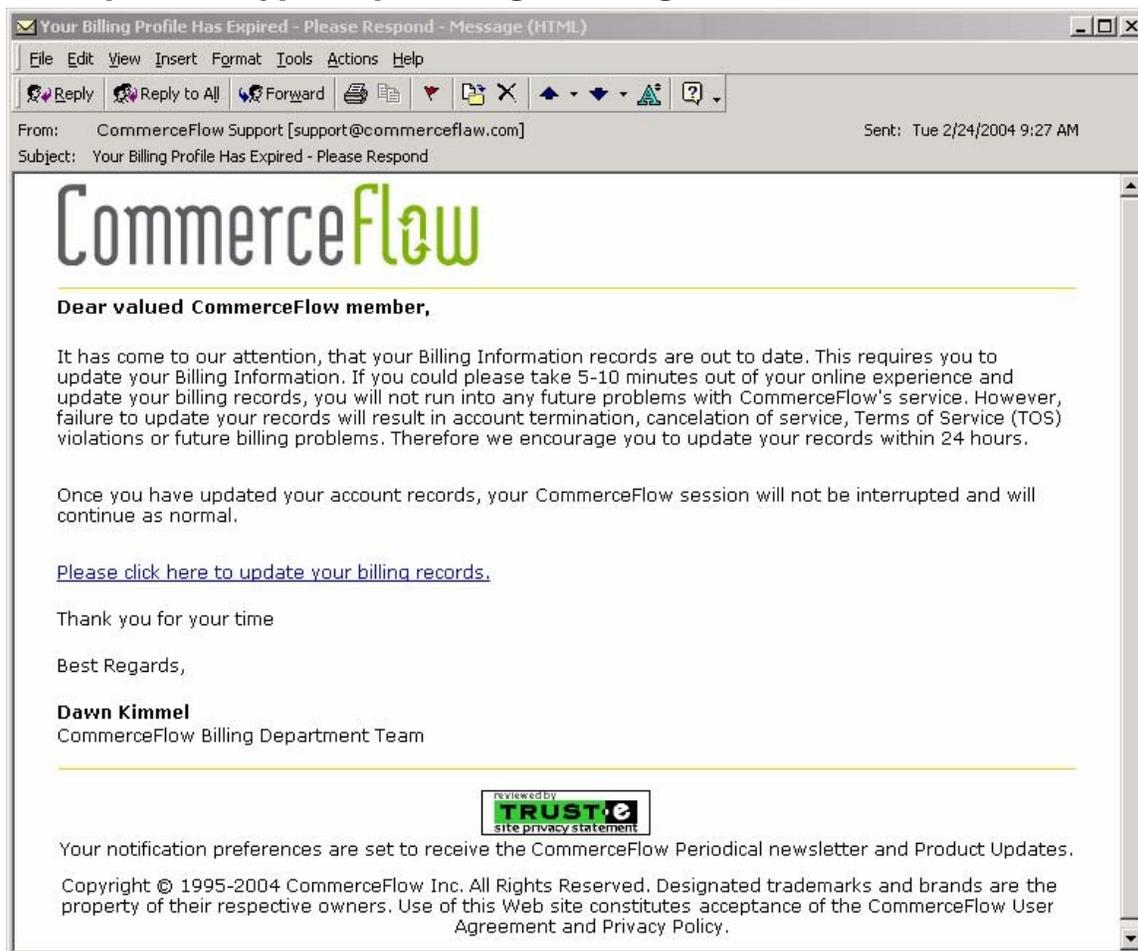
While browser technologies are constantly being updated to address recent phishing tactics, browsers are large, complex programs that must provide considerable functionality and flexibility to satisfy the needs of legitimate web site designers. It is highly improbable that deceptive phishing appearances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

Deceptive information flow

To maximize the value of a compromise, the user should not know that he or she has provided confidential information to a phisher. After obtaining confidential information, phishing sites often inform the user that he or she must log back into her account now that the information is “confirmed,” and redirect to the legitimate site.

Phishers sometimes construct elaborate information flows to cover their tracks and conceal the ultimate destination of compromised information. In some cases, these information flows can contain multiple media, such as compromised “zombie” machines, instant messaging, and anonymous peer-to-peer data transfer mechanisms.

Example of a typical phishing message

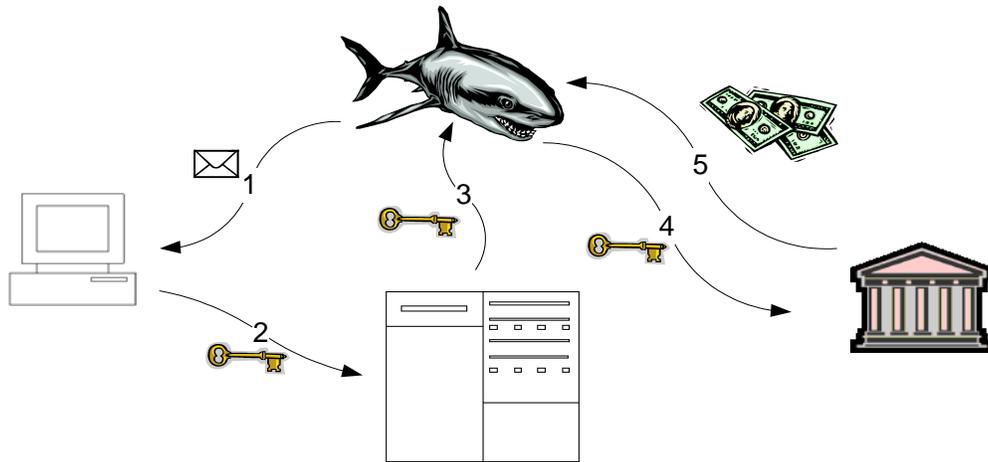


This particular phishing message was sent using a deceptive return address ("commerceflow.com").* It contains deceptive imagery from both CommerceFlow and Trust-e, presents a fraudulent call to action on the pretext that account information is outdated, and cloaks the destination of a clickable link to trick the user into believing the link refers to the genuine CommerceFlow site.

Technology Countermeasures

To best understand the context in which phishing countermeasures operate, it is important to understand the information flow in a phishing attack.

* Innumerable companies have been targeted by phishing attacks. To avoid singling out any particular victim, this document uses the generic companies "CommerceFlow" and "Large Bank and Trust" to represent legitimate, customer-trusted institutions.



The simplified flow of information in a phishing attack is:

1. A deceptive message is sent from the phisher to the user.
2. A user provides confidential information to a phishing server (normally after some interaction with the server).
3. The phisher obtains the confidential information from the server.
4. The confidential information is used to impersonate the user.
5. The phisher obtains illicit monetary gain.

Steps 3 and 5 are of interest primarily to law enforcement personnel to identify and prosecute phishers. The discussion of technology countermeasures will center on ways to disrupt steps 1, 2 and 4, as well as related technologies outside the information flow proper.

Preventing a phishing attack before it begins

Before steps 1-5 above, a phisher must set up a domain to receive phishing data. Pre-emptive domain registration may reduce the availability of deceptively named domains. Additionally, 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 deceptively named domains, but would not address the ability of phishers to impersonate sites. As email authentication technologies become more widespread, email authentication could become a valuable preventive measure by preventing forged or misleading email return addresses.

Some services attempt to search the web and identify new phishing sites before they go “live,” but phishing sites may not be accessible to search spiders, and do not need to be up for long, as most of the revenues are gained in the earliest period of operation. The average phishing site stays active no more than 54 hours.

Detecting a phishing attack

Many different technologies may be employed to detect a phishing attack, including:

- 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.
- Monitoring call volumes and the nature of questions to customer service.
- 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 will 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 for the attack in a timely manner.

Preventing the delivery of phishing messages

Once a phishing attack is underway, the first opportunity to prevent a phishing attack is to prevent a phishing message from ever reaching a user. This represents a disruption of step 1 of the phishing information flow.

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 a user. 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.

Phishers depend on being able to make their messages visually appear to be from a trusted sender. One possible countermeasure is to detect unauthorized imagery 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. This means that 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.

Authentication

Message authentication techniques such as Sender-ID have considerable promise for anti-phishing applications. Sender-ID prevents return address forgery by checking DNS records to determine whether the IP address of a transmitting mail transfer agent is authorized to send a message from the sender's domain. Yahoo! Domain Keys provides similar authentication, using a domain-level cryptographic signature that can be verified through DNS records. Some form of lightweight message authentication may be very valuable in the future in combating phishing. For the potential value to be realized, Sender-ID or a similar technology must become sufficiently widespread that invalid messages can be summarily deleted or otherwise treated prejudicially, and security issues surrounding the use of mail forwarders need to be resolved.

In the future, it is possible that cryptographic signing (such as using S/MIME) may become sufficiently prevalent that unsigned emails may be discarded or treated prejudicially. As unsigned emails constitute the vast majority of messages today, it is presently impractical to institute such a practice.

Preventing deception in phishing messages and sites

There are two different points to thwart phishing presentation deception: at the message, and at the site to which the message points. In the overall diagram of the phishing information flow, both of these countermeasures prevent step 2. Since these two choke points require similar technologies, they are discussed together.

Signing

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.

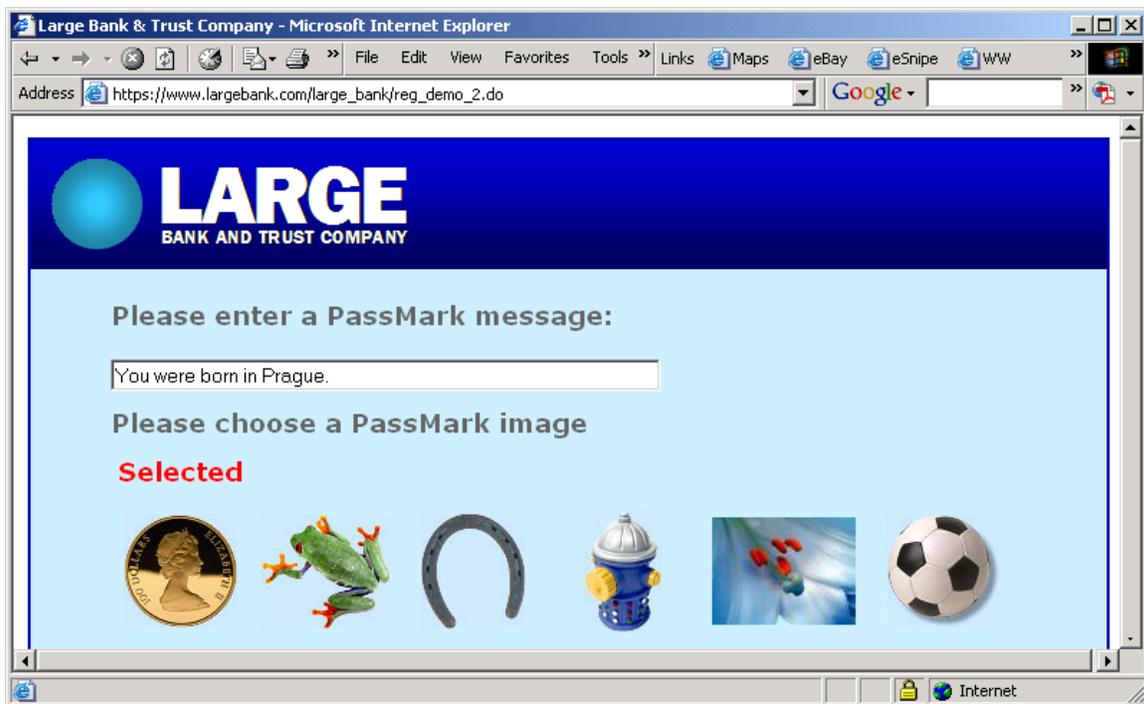
Personally identifiable information

The simplest way to reduce the deceptiveness of phishing messages is to include personally identifiable information with all legitimate communications. For example, if every email from bank.com begins with the user's name, and

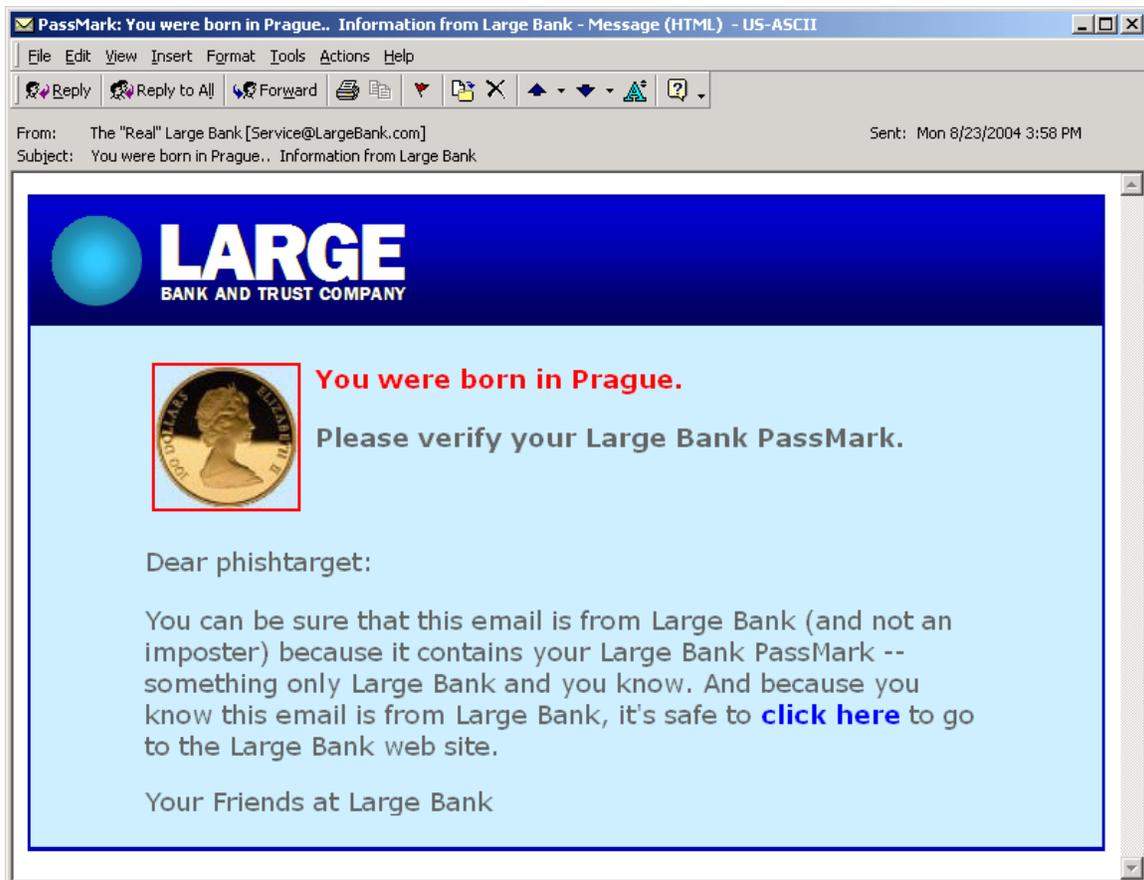
every email from bank.com educates the user about this practice, then an email that does not include a user's name is suspect. While implementing this practice can be complex due to the widespread use of third-party mailing services, it is an effective measure.

Beyond static identifying information, more sophisticated personally identifiable information may be included, such as text that a user has requested to be used. This would permit 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.



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



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 is legitimate before displaying personally identifiable 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 mean that a user is more likely to discern differences between a phishing attack and a legitimate interaction.

Indication of suspicious content: relying on the user

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, so a user, however well-educated, cannot be relied on to discern between a legitimate message and a phishing attack.

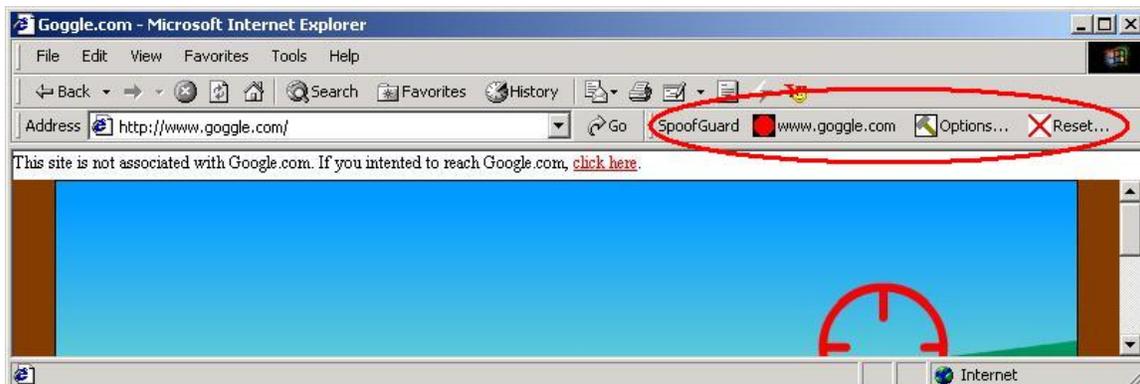
Even for criteria that may not be readily spoofable, passive measures that rely on the user to look for a warning sign will have questionable efficacy. It is worth educating users, but education cannot be relied on as a primary defense.

One form of providing an explicit indication of unsafe content is to increase information sharing, which is discussed in more detail below. If an email was deemed to be questionable by an anti-spam or anti-phishing content filter, but not poor enough to be undeliverable, it could be highlighted to indicate that the user should be wary.

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:



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:



Anti-phishing toolbars generally combine a visual safety indication with outbound data monitoring to attempt to prevent disclosure of confidential information to unauthorized parties, as discussed below.

Vendor-specific anti-phishing toolbars are a good preventive measure. Many users use multiple services that could benefit from such protection, and it is not practical to install a separate toolbar for each one. In the long term, it will be

necessary to combine knowledge about multiple sites into a single unified toolbar.

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.

Canonical display of deceptive content

Presently, a web designer may design a page with links to be displayed however he or she wants. This makes it easy to create phishing pages. 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 as suspicious to the user. For example, consider the following HTML fragment:

```
<CENTER><H1>Suspicious URLs</H1></center>
<P>To go to a surprising place via a cloaked URL, click on
<A HREF="http://security.commerceflow.com@phisher.com">this link.</A>
<P>To go to a surprising place via a cloaked URL with a password, click on
<A HREF="http://security.commerceflow.com:password@phisher.com">this
link.</A>
<P>To go to a surprising place via an open redirect, click on
<A HREF="http://redirect.commerceflowsecurity.com?url=phisher.com">this
link.</A>
<P>To go to a surprising place via misleading link, click on
<A HREF="http://phisher.com">http://security.commerceflow.com.</A>
```

This will normally render as:



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. A 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 Javascripts to modify the status bar when a URL is being shown). The page above might be rendered as:



Modifying images on the fly

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.

Interfering with the call to action

A phishing attack using email and a browser asks a user to perform an action, such as clicking on a link. One class of countermeasures focuses on disrupting the initial call to action. This class of countermeasures is another form of defense against step 2 of the phishing information flow.

Increasing information 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 a spam filter, an email client and a browser. A spam filter may have classified a message as being possible spam, 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 Microsoft.

Information gleaned while processing messages can help thwart phishing. If an email is known to be 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. Scripts can be disallowed, links can be shown with their true names, forms can be disallowed, etc.

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, capabilities (scripting, form submissions, display of links, etc.) can be restricted for links pointed to in less trustworthy messages.

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

Warning about unsafe actions

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 could be usefully included:



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.

Interfering with transmission of confidential information

Another point at which phishing attacks may be disrupted is when a user attempts to transmit confidential information (step 2 of the phishing information flow). 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.

Outgoing data monitoring

One class of technology to intercept the transmission of confidential information is the toolbar approach. A browser plug-in such as a toolbar can store hashes of confidential information, and monitor 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. This approach has a challenging obstacle to overcome. Phishers may scramble outgoing information before transmitting it, so keystrokes must be intercepted at a very low level. 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 presently most phishing attacks do not include effective countermeasures.

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 2 of the phishing information flow. However, 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.

Similarly, a blacklist based on hostnames is also 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.

Domain-specific passwords and password hashing

Phishing for passwords only works if the password sent to the phishing site 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. A proposed implementation of this idea is called password hashing. This method hashes password information with the domain name to which it is going, so that the actual transmitted passwords can be used only at the domain receiving the password data. Such hashing could be provided by a browser as a built-in mechanism that is automatically performed for password fields. This provides excellent data security for compromised sites as long as passwords are difficult to guess through a dictionary attack, in that stolen password data cannot be applied to any other site. However, the user still types in his or her usual password in a browser to gain account access, and it would be difficult to prevent phishers from simulating password input, bypassing any hashing, to capture the raw password data. If combined with reserved screen real estate for password entry, password hashing would be rendered less susceptible to attack. Password hashing, as with the secure data transmission discussed below, is aimed at making data gleaned in step 2 of the phishing information flow unusable in step 4.

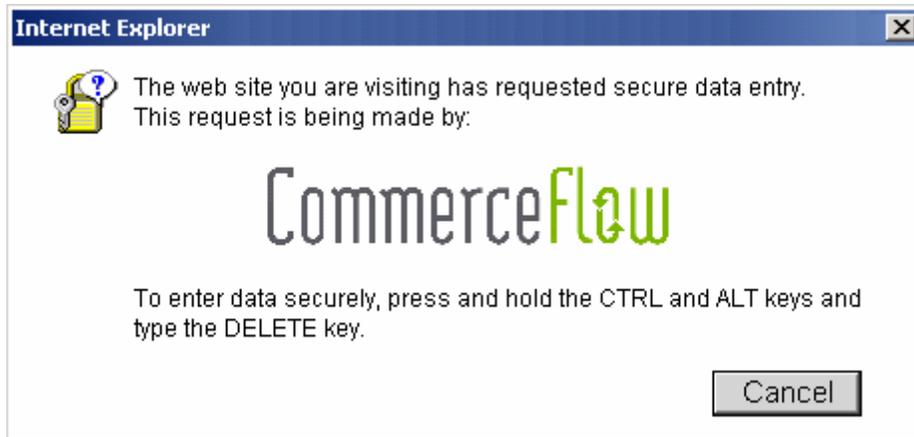
Secure data transmission

A more radical future-focused technology to render information useless to phishers is to provide a non-spoofable secure data entry service at the operating system level, accessible from a browser, ensuring that sensitive information can reach only a legitimate recipient. This has been done at the operating system level for login information through 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 Windows computer, which was implemented as part of the National Computer Security Center's requirements for C2 certification. Today, the threat of data entry theft through phishing is more heterogeneous than in the past, and extensions to this tested mechanism could be applied to thwart phishing attempts.

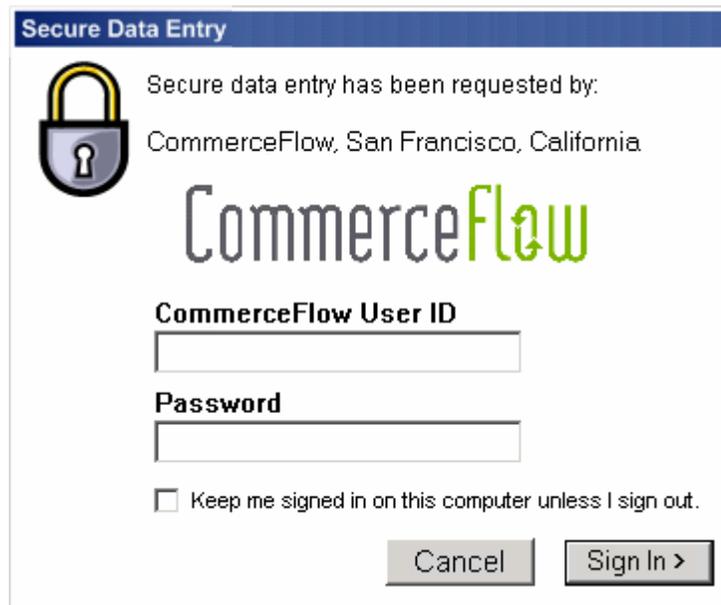
An operating system could safeguard the entry of sensitive information by providing a secure data entry service that is called with two separate types of arguments:

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

After the operating system has been notified of the impending trusted data entry, the user is prompted to hit CTRL-ALT-DEL, a non-interceptable key sequence under Windows which transfers control to the operating system.



When the user enters the secure key sequence, the operating system determines that secure 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.



This 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, absent a root 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 who possesses the corresponding private key can read the data. This encrypted data is then made available to the requesting application. If a certificate is presented by someone who doesn't own it, they will be unable to interpret the sensitive data, as only the certificate owner has the private key needed to decrypt it.

This particular secure data entry mechanism relies on certification authorities to verify the identity and logo of an applicant before granting a certificate. Secure

data entry certificates would be issued by a small, controlled set of authorities, who can be trusted never to issue a certificate without proof of identity and to ensure that an unauthorized logo is not being used. Requirements to be a trusted certification authority for secure data entry should be at least as stringent as the requirements for root certification authorities for SSL certificates, and possibly more stringent.

Unlike ineffective admonitions to check an advisory display element such as the lock icon, getting to a trusted path data entry screen is an active part of the user experience. As users grow accustomed to always entering sensitive data (passwords, credit card numbers, social security numbers, etc.) through a trusted path mechanism such as typing CTRL-ALT-DEL, any request for sensitive data through an insecure page or message would raise an immediate red flag – which could be augmented by a detection system indicating data transmission to an untrusted site, or entry of sensitive data.

Interfering with the use of compromised information

Another technology-based approach to combating phishing is to render compromised information less valuable. Apart from technologies to render information irretrievable, such as hashing passwords with domains and a trusted path that encrypts information with a public key, additional requirements may be placed on the use of information to mitigate the impact of compromise. These technologies attack step 4 of the phishing information flow.

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)
- What you know (e.g. an account name and password)

Phishing attacks typically compromise what a user *knows*. In a remote computing environment such as the internet, it is difficult to ascertain what the user *is*, so the usual second factor is to verify something that the user *has* in addition to account information. In order for this to be effective, two-factor authentication must be required for every transaction. For example, a user must have a USB dongle, or type in a time-sensitive code from a hardware device, or swipe a smart card. This is a highly effective measure, though expensive in the cost of purchasing and distributing security devices, the deployment of infrastructure for reading them, and the inconvenience to customers in using them. Conventional two-factor authentication is appropriate for high-value targets such as commercial banking accounts, but so far has not taken root in the United States for typical consumer applications.

Light-weight two-factor authentication

A less costly approach to two-factor authentication is to have a device identifier, such as a checksum of all available machine information, which can authenticate the device. Such a device identifier must be transmitted only to a secure location, or employ other measures to prevent man-in-the-middle attacks. This has the advantage of not requiring additional hardware, and the disadvantage that it does not permit a user to use normal transaction authorization procedures when away from an authorized machine.

Another problem: Cross-site scripting

Cross-site scripting is an alternative to step 1 in the overall phishing information flow, in which rather than sending an email, a phisher inserts malicious code into a web page of a targeted institution. Any web page that contains externally supplied information, such as an auction listing, product review or web-based email message, may be the target of a cross-site scripting attack. Once inserted, a 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 (a combination of steps 2 and 3 in the phishing information flow).

Filtering out cross-site scripting

Any user data that is ever displayed on the screen should be filtered for cross-site scripting. 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.

Browser security enhancements to prevent cross-site scripting

There are many ways in which cross-site scripting may 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 cross-site script from including a valid `</noscript>` tag and inserting cross-site scripting, a dynamically generated random key should be used that must match in the `<noscript>` and `</noscript>` tags. 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]

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.
- Provide Sender-ID information in DNS records for your mail servers. 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 personally identifiable information in each email to a customer. Along with the personally identifiable information, include an educational statement that it is your policy always to do so.
- Provide an email address such as spoofer@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.
- 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 loans) 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.

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 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.
- Personally identifiable 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 distinctive display of potentially deceptive content and providing a warning when a potentially unsafe link is selected, could substantially reduce the efficacy of phishing attacks.
- 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.
- Anti-phishing toolbars are promising tools for identifying phishing sites and heightening security when a potential phishing site is detected.
- Detection of outgoing confidential information, including password hashing, is a promising area of future work, with some technical challenges.
- An OS-level trusted path for secure data entry and transmission has the potential to dramatically reduce leakage of confidential data to unauthorized parties.
- Two-factor authentication is highly effective against phishing, and is recommended in situations in which a small number of users are involved with a high-value target. Device identifier based two-factor authentication offers the potential for cost savings.
- Cross-site scripting is a major vulnerability. All user content should be filtered using a let-in filter. Browser security enhancements could decrease the likelihood of cross-site scripting attacks.

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 Secret Service 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. Examples of vendors with solutions in this category are:

- OSpam.net
- 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 Personally Identifiable Information

- PassMark Security

Email Filtering to Remove Fraudulent Email

Companies in this category provide software that is run on the customer premises that can remove many phishing emails as well as spam. Vendors include:

- OSpam
- Brightmail
- Engate
- Ironport
- MailFrontier

Identifying a Valid Web Site

Identifying a valid web site, or warning about a potentially fraudulent web site.

- Billeo
- PassMark Security
- Stanford SpoofGuard
- Whole Security

Providing stronger authentication

Two-Factor Authentication

- Software PKI Certificates
 - GeoTrust
 - Thawte
 - Verisign
- Key Fobs
 - RSA SecureID
 - SafeNet iKey
 - Secure Computing
 - Thales
- Smart Cards
 - ActivCard
 - Gemplus

- Virtual Second-Factor Authentication
 - Anakam
 - Arcot Systems
 - PassMark Security
- Biometric
 - Bio-Key International
 - Bioscrypt
 - DigitalPersona

Desktop Technologies

Toolbars and Phishing Site Detection

- Billeo
- GeoTrust
- Stanford SpoofGuard
- Whole Security

Malware Detection

- McAfee
- Sophos
- Symantec
- WebRoot
- WebSense
- Whole Security

Consulting Services

Education

- Glennbrook Partners
- Internet Identity

Security Technology Evaluation and Development

- Radix Labs

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.ussstreas.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 gather 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

- Steven M. Bellovin, *Using the Domain Name System for System Break-ins*. Proceedings of Fifth Usenix UNIX Security Symposium, June 1995.
- 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.
- 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.
- 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.
- 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).
- Robert T. Morris, *A Weakness in the 4.2BSD UNIX TCP/IP Software*. Computing Science Technical Report 117, AT&T Bell Laboratories, February 1985.
- Rod Rasmussen, *Phishing Prevention: Making Yourself a Hard Target*. Internet Identity / APWG (April 5, 2004).
- Blake Ross, Nick Miyake, Robert Ledesma, Dan Boneh and John C. Mitchell, *A Simple Solution to the Unique Password Problem*.
- Z. Ye. *Building Trusted Paths for Web Browsers*. Master's Thesis. Department of Computer Science, Dartmouth College. May 2002
- 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

Appendix E

The San Francisco Electronic Crimes Task Force

This document was prepared by members of the San Francisco Electronic Crimes Task Force. This task force 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 task force are:

Val Batiste	Wells Fargo
Ken Beer	Tumbleweed Communications
Dan Boneh	Stanford University
Frank Christian	United States Secret Service
Drew Dean	SRI International
Aaron Emigh	Radix Labs
Louie Gasparini	Passmark Security
Rajesh Lalwani	Billeo
Karl Levitt	University of California, Davis
Tom Lickiss	United States Secret Service
Patrick Lincoln	SRI International
Dan Maier	Tumbleweed Communications
John Mitchell	Stanford University
Joyce Reitman	emsPartners
Jim Roskind	Radix Labs
Jeff Rowe	University of California, Davis
Abe Smith	Xilinx
Doug Tygar	University of California, Berkeley
Don Wilborn	United States Secret Service

Communications regarding this document may be directed to Aaron Emigh, ate@radixlabs.com.