

Anti-Phishing Best Practices for ISPs and Mailbox Providers: A MAAWG Document.

Vipul Ved Prakash, Founder and Chief Scientist, Cloudmark, Inc.

INTRODUCTION

The entire Internet community is familiar with spam and phishing attacks. While spam is an annoyance, phishing can cause major financial disruptions for those involved. The phishing problem is becoming a major concern for the ISP market, and the pressure is coming from both users who are beginning to demand that their service providers take action against the attacks and from financial institutions that are the target of the attacks.

ISPs are being forced to actively participate in the global reduction of phishing attempts to mitigate customer churn and possibility of litigation. There is a lack of consensus as to how an ISP should attempt to reduce the exposure of its users and the global community to phishing attempts. This document distills the best of the practices used by members of MAAWG to combat phishing attacks.

Nature of phishing attacks

The Internet community and ISPs have a reasonable handle on the spam problem today. In-house technology, third party solutions and industry initiatives have focused on anti-spam for a long time and most ISPs are equipped to combat spam with acceptable efficiency. Phishing is a relatively new threat, a more insidious one at that, which is further exacerbated by its surface similarity to spam. Phishing attacks are relatively sophisticated and logistically different from spam and understanding the differences is critical when considering a framework for protection from these attacks. This section examines the more important differences between spam and phishing.

* Sophisticated

The first important difference about Phishing is that it is relatively more sophisticated when compared to Spam. Not only are the messages that Phishers send out carefully crafted to impersonate known, trustworthy financial institutions or organizations, but Phishers also systematically exploit software vulnerabilities in web browsers, web servers and local operating systems in order to fool filtering software, trick users, and steal as much information as possible. Their methods and techniques bear a much closer resemblance to blackhat hackers than to Spammers.

* Targeted

Unlike Spammers who target large groups, lists, and user bases in an attempt to capitalize on some given advertising response rate, Phishers instead target highly qualified lists of email addresses and individuals not just to improve response rates, but more importantly to evade less sophisticated data collection (e.g. honeypots) and detection systems which might alert authorities to the occurrence of their fraud attempt.

* Transient

Phishing attacks are also very transient and short-lived, often occurring for only a few hours before disappearing, which is in contrast to Spam which tends to be sent out in frequent and large batches. Since Phishing is a well-defined criminal activity, with constraints that are quite different from Spam, transient attacks become very necessary to the phisher in order to evade detection.

* Dynamic

Finally, Phishing attacks and the sites that host them are very dynamic - they move between servers very quickly. Whereas a Spammer is advertising some product or service from a known website, a Phisher is redirecting users to a private webserver/site specifically designed to

impersonate some financial institution or organization. Phishers typically exploit software vulnerabilities in web servers or server operating systems in order to install their own content, and because of their increased sophistication Phishers are able to do this in an automated fashion, allowing them to compensate quickly and easily when a compromised site is discovered and taken offline. In fact, Phishers cycle through compromised hosts quickly regardless of their discovery, simply to confuse or obfuscate the true source of the attack.

Inbound protection

In bound filtration of phishing messages

Inbound mail filtration of phishing emails through the use of anti-phishing/anti-spam filtration technology either at the end-point or MTA is one of the front-line defenses against phishing emails. This is carried out in much the same way, and usually through the same or similar software, as anti-spam filtration.

Several techniques have been developed and are in use for filtration of spam. These include, IP based blacklists, Bayesian filters, heuristics engines, content fingerprinting schemes, and sender authentication. While all these techniques are effective to varying degrees against spam, only some perform well against phishing. Here's a breakdown of what to expect with different techniques:

Bayesian Filter

Bayesian classifiers filter spam based on their semantic difference from legitimate communications. Bayesian filters are in wide-use, especially in end-point anti-spam products, where they perform best. While the content of spam differs considerably from the content of legitimate communications, phishing messages – that imitate legit content - tend to look like legit communications and are classified as such by Bayesian classifiers trained to detect spam. It is rather hard for vanilla Bayesian classifiers to distinguish between spam and phishing. However, a Bayesian classifier specifically trained to detect phishing messages can do a good job. ISPs that deploy Bayesian filters should carefully measure the effectiveness of their filters for phishing messages.

IP Based Blacklists

The previous section describes how phishers use compromised machines to host phishing web pages and to send out phishing emails. IP or source based filters, created to address spam, are particularly poor at detecting phishing messages because such messages often originate from "good" hosts. As with Bayesian filters, ISPs should evaluate the quality of their source based blacklist solutions for filtration of phishing messages.

Heuristics and Fingerprinting schemes

Heuristics and fingerprinting schemes tend to perform best against phishing, specially if such solutions are targeted against phishing attacks. Both heuristic and fingerprinting systems tend to rely on the specific nature of phishing attacks and do not have the shortcomings of Bayesian and source-based filtration.

URL based filters

Several URL based filters look for specific IPs, domains or URLs where phishing web pages are hosted. These IPs, domains and URLs are collected from reports of phishing messages gathered from email users and honeypots. Such filters are fairly effective, however, since phishers tend to use a large set of hosts and cycle hosts frequently, it is hard to have a comprehensive and updating list of bad IPs, URLs and domains based on limited reporting.

For more information visit us at
www.cloudmark.com

Headquarters

128 King Street, 2nd Floor
San Francisco, CA 94107 USA
Ph: +1.415.946-3800
Fax: +1.415.946-3871

Cloudmark Europe Ltd.
Garrick House
26-27 Southampton Street
London
WC2E 7RS
England
Ph: +44 (0) 207-100-5224

Asia Pacific Office
45/F The Lee Gardens
33 Hysan Avenue
Causeway Bay
Hong Kong
Ph: +852 3180 7768

Recommendations:

1. The ISP should evaluate inbound phishing filtration solutions based on the inherent strength of the underlying technology.
2. When considering a third-party solution, the ISP should look for solutions with explicit anti-phishing support.
3. The ISP should conduct comprehensive tests before deploying an anti-phishing solution.

Policy considerations

Spam is often tagged and allowed to go through for users to make the final decision. With Phishing messages, it is advised that the ISP drop the message (or reject it at SMTP level), because such messages tend to trick users into thinking that they are legitimate and the ISPs filtration software made an incorrect decision.

Recommendations:

1. Drop phishing messages where possible.
2. When dropping messages is not possible (due to user request, ISP policy or legislative requirements), the ISPs should tag messages to indicate that they are phishing messages and while they might look legitimate, they are dangerous and should be ignored.

End-point filtration

There are several free and pay-for end-point security solutions on the market that plug-in to user's email software and filter phishing messages from incoming mail. These solutions can be effective when the ISP is unable to provide server level filtration of phishing email. End-point solutions are also recommended by ISPs so their users can be protected when they are accessing email from multiple accounts, some of which do not reside on the ISP infrastructure. Also, end-point security solutions are invoked when a user reads their email, as opposed to server-side solutions that are invoked when the mail is delivered. Often, the latency between delivery and processing of mail is long enough for end-point filters to be updated and hence provide better security. The latency is a good argument for providing two-tier protection from phishing.

Recommendations:

1. ISPs should encourage their users to employ end-point security solutions to combat phishing. Forgery detection with Sender Authentication Email authentication standards, SPF, Sender-Id and DKIM, are becoming widely adopted. Among other things, email authentication can be used to determine if the sender has forged the sender identity. Phishers often try to forge the information in the headers to make it appear as if originating from a legitimate institution. Sender authentication, where available, can be used to detect this.

Recommendations:

1. ISPs should filter or drop email if forgery can be unequivocally established.

Visual Identification

In parallel to the continued effort related to blocking phishing messages, some leading mailbox providers will support a mechanism to convey authenticity of good messages to their users. AOL and Yahoo! announced they indicate legitimacy of good email to the recipient using a visual cue. The visual cues should appear in non-spoofable UI areas. Another proposed mechanism is to provide folders that only contain trustworthy email, so all other email (including phishing messages) is automatically considered suspect

Recommendations:

1. ISPs should explore visual authentication and identification technologies for legitimate email. Hide images from untrusted sources. Displaying images included in untrusted email messages puts recipients at great risks. Security vulnerabilities such as the recently publicized WMF exploit[x] underline the need to protect users not only from subjectively offensive images (i.e. pornography) but also from images that could abuse security breaches and install key loggers and other malware on machines of unsuspecting users. Email providers have long disabled JavaScript and other executables for all incoming email messages. There is now a positive trend (AOL, Gmail, Outlook 2003, Window Medial Live) to disable images by default and to display images only when embedded in trustworthy messages.

Recommendations:

1. ISPs should consider turning off images for all messages for which the identify and reputation of the sender cannot be established. Remove links from phishing emails. Another method of alleviating the threat of phishing is to replace links to phishing sites present in phishing email with links that point to educational sites about the threat of phishing. If the ISP is forced to deliver phishing messages to the user, link removal can be a big help.

Recommendations:

1. ISPs should remove links to phishing sites from phishing emails if they have to deliver the phishing messages to end-users.

Web Traffic Filtration

Phishing messages contain one or more links to the phishing website that collects user credentials. One way to render phishing attacks useless is to block access to these sites. There are several open and commercial efforts underway that list phishing URLs and provide these lists to organizations that wish to limit access these URLs. [Q: Should we provide a list of organizations that build phishing URL lists?]

There are several tools that empower the end-user with knowledge about the authenticity or fraudulence of website they visit in their web-browser can effectively curb phishing victimization even if users click on links within phishing emails. Such tools are designed to "plug-in" to the web browser and they examine the links as well as the contents of the visited web pages to make a decision regarding the safety of visited web-pages. Web-browser anti-phishing toolbars add an additional last line of defense.

Recommendations:

1. Where possible, ISPs should enable short-lived blocks on confirmed phishing sites using firewalls and/or web-filtration products.

2. ISPs should bundle, distribute or encourage their users to download web browser plug-ins that detect and restrict access to phishing sites based on phishing URL feeds and/or predictive heuristic technologies. Certain plug-ins also authenticate good websites and instill confidence in users about safety of their web experience.

OUTBOUND PROTECTION

Phishers launch their attacks from compromised hosts. Phishers will either generate the phishing e-mails directly from their server or redirect the messages through a botnet, or a large array of compromised nodes under their control. In either case, the malicious traffic is transported by an unsuspecting carrier. Often, these unsuspecting carriers are customer machines hosted on ISP infrastructure. Phishers will use the ISP email infrastructure to send out phishing emails.

It is within the capabilities of an ISP to make an attempt at filtering outbound phishing attempts using anti-phishing filters. Several third-party filtration solutions operate in "out-bound mode" and can be used to stop phishing messages from leaving the ISP network.

The other advantage of an out-bound filter is that it might provide the ISP with a report indicating the location of phishing web-pages. If these web-pages are installed inside the ISP infrastructure, the ISP can take a decision to remove them or restrict access to them.

Recommendations:

1. The ISP should consider third-party out-bound phishing filters. When considering an in-bound filter, the ISP should also evaluate the out-bound capabilities of the solution.

Pharming and DNS cache-poisoning protection

Phishing attacks normally arrive via email and request that a user visit a counterfeit site to enter personal information. Links to the phony site are always provided, and by looking at the URL address, it's possible to determine if the site is indeed a fake. However, the successor to phishing attacks, Pharming, can not be as easily spotted. While a user may believe that they are visiting ebay.com, they may be unknowingly visiting a counterfeit site.

Pharming attacks work by attacking, or poisoning, the DNS systems used to translate Internet addresses. While attacks of this nature have been around for years, they have just recently come into light as method for identity theft attacks. A Pharmer with access to a DNS server will isolate specific sites, and route Internet traffic from the real site to a scam site. Unbeknownst to the user, he or she is now visiting a phony site, and any information entered could be used for malicious reasons. The unique danger of pharming attacks, as opposed to phishing attacks, is that a phony link does not need to be provided; even typing in a valid address by hand can still lead to an attacker's site. Also, individual users do not need to be singled out, as any amount of users who use the compromised DNS server are now vulnerable to an attack.

Recommendations:

1. Currently, the best practice for avoiding pharming scams is to only surf sites with certificates enabled. The ISP should educate their users about the dangers of pharming and encourage them to check for certificate validation when providing important credential information to a web-site.
2. ISPs should deploy DNSSEC (short for DNS Security Extensions), which provides digitally signed authentication data in order to validate the site in which the user is visiting. While DNSSEC will not prevent an attack from happening, it will inform the user of DNS data that is untrustworthy.
3. ISPs should ensure their DNS architecture is up-to-date. Old software or systems may be vulnerable to attacks, which could lead to the compromise of a DNS server, and put all users of that server at risk for pharming attacks.

Phishing related customer support calls

Phishing problems inevitably generate support calls and effective processes that support can interface with to streamline the remediation of the support tickets save valuable time. A phishing or pharming support call is also a great vector for user education.

Recommendations:

1. If a user reports suspicious email asking for personal information, the ISP should inform the user of the dangers of phishing attacks, and warn him or her against giving out personal information online. The user should be further advised to send a copy of the email to the ISP, so it can be used to build better filters, and copy both FTC and APWG.
2. If the user believes that he or she has been scammed, suggest that the user file a complaint with the FTC at www.ftc.gov, and visit the FTC's Identity Theft Web Site at www.consumer.gov/idtheft.
3. If the suspected phishing email or site is hosted or was sent from the ISP, there should be in place a policy and process that support can interface with to remediate the problem site or computer.
4. The customer support should also direct the user to a consumer education resource that enables them to understand the nature and scope of the threat, as well as the measures in place (at the ISP) for protecting users from phishing.

ISP to Financial institution communications

ISPs should try to communicate early knowledge of phishing attacks to the financial institution at which the attack is directed.

Recommendations:

1. The ISP should communicate knowledge of phishing attacks to the targeted financial institution via an organization the Anti Phishing Working Group or a similar regional organization.

For more information visit us at
www.cloudmark.com

Headquarters

128 King Street, 2nd Floor
San Francisco, CA 94107 USA
Ph: +1.415.543.1220
Fax: +1.415.543.1233

Cloudmark Europe, Ltd.
Carmelite, 50 Victoria Embankment
Blackfriars, London EC4Y 0DX UK
Ph: +44 (0)207.100.5224
Fax: +44 (0)207.100.5224