

PayPal[™]

A Practical Approach to **Managing Phishing**

Michael Barrett, Chief Information Security Officer
Dan Levy, Senior Director of Risk Management – Europe

April 2008

Index

1.0 Executive Summary

- 1.1 Introduction
- 1.2 PayPal's View of Phishing
- 1.3 Our Response

2.0 Definitions, Scope and Principles

3.0 Reclaiming Email

- 3.1 Email Signing and Blocking
- 3.2 Visual Identification

4.0 Block Phishing Sites

- 4.1 Unsafe Browsers
- 4.2 Blacklists
- 4.3 Anti-fraud Warning Pages
- 4.4 Extended Validation SSL Certificates

5.0 Ancillary Strategies

- 5.1 Customer Education
- 5.2 Site Shutdown
- 5.3 Authentication
- 5.4 Fraud Models
- 5.5 Law Enforcement
- 5.6 Government Relations

6.0 Results

7.0 Conclusions and Next Steps

1.0 Executive Summary

1.1 Introduction

Surely only a few individuals, who have been living a life of seclusion on the French Riviera for the last few years, won't know what the crime of "phishing" is. At least that's what we tend to think in the security industry. Yet, according to Gartner estimates, 3.3% of the 124 million consumers who received phishing email last year were victimized and lost money because of the attacks. In short, phishing is a "con trick" by which consumers are sent email purporting to originate from legitimate services like banks or other financial institutions. The email, which generally contains a call to action such as "update your account details," contains a link to a website where the consumer is asked to provide their log-in credentials for the legitimate site, often along with other personal and confidential information such as a bank account number, credit card number, social security number, or mother's maiden name. Once the information has been collected from the legitimate owner, it is then used to conduct various kinds of fraud, including identity theft.

In the summer of 2006, the authors of this white paper examined PayPal's approach to managing phishing. We realized that our strategy was based on preventing financial loss in the victim's account – long after the original phishing email had duped its victim. However, it became rapidly clear to us that there was a holistic dimension that our previous approach missed. Equally clear was the fact that we couldn't eradicate this problem on our own – to make a dent in phishing, it would take collaboration with the Internet industry, law enforcement, and government around the world.

1.2 PayPal's View of Phishing

Like any successful businessperson, fraudsters are driven by profit. Their profit equation is a simple formula:

$$V * R * M = P$$

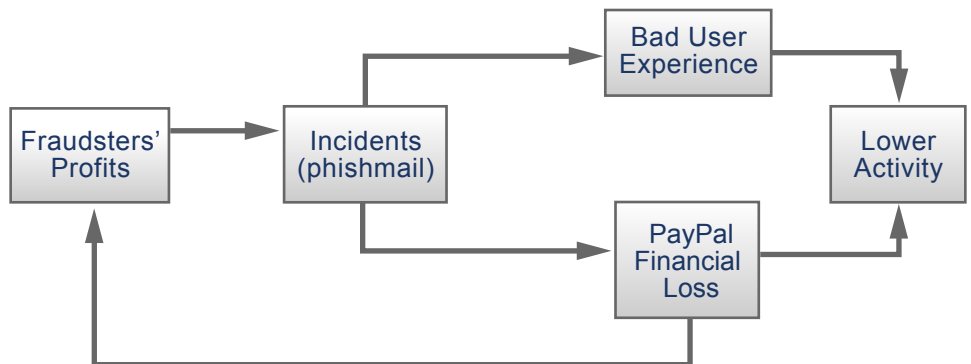
- V = Volume of phishing sent by fraudster**
- R = Response rate (percent of victims that give up account information)**
- M = Monetized value of a stolen account**
- P = Profit**

In the past, PayPal has focused successfully on the third input – the 'monetization' of the stolen accounts. This had the dual effect of initially lowering the fraudsters' profits (and lowering PayPal's losses), but also encouraging the fraudsters to increase the phishing sent (V) in order to maintain or grow the absolute level of their profits.

While our efforts to reduce losses to PayPal and our sellers worked, fraudsters simply increased their phishing volume. By some industry calculations, phishing purporting to be from PayPal and eBay reached more than 75% of total phishing sent.¹

We therefore reassessed the situation and drew the picture you see in Diagram 1: *PayPal's View of the Phishing Problem*, on our whiteboard.

Diagram 1: PayPal's View of the Phishing Problem



¹"Over 75% of all phishing emails target PayPal and eBay users," Sophos, July 2006

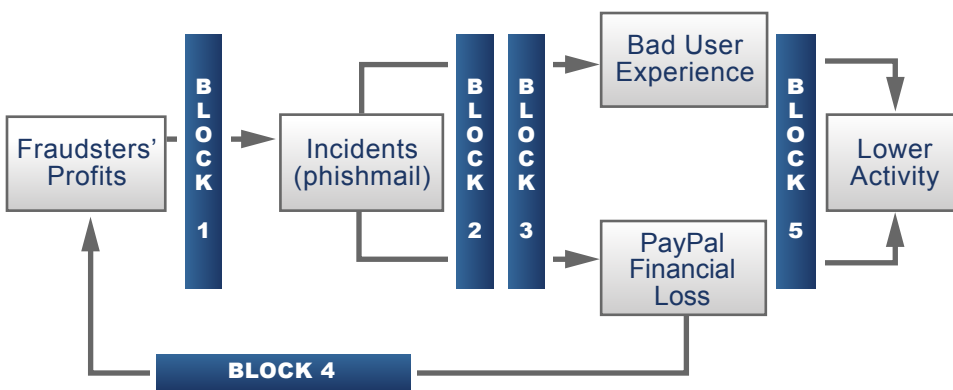
We knew that fraudsters' profits drove them to send phishing mail. In turn, phishing mail drove two responses from our customers: a bad user experience and/or financial loss. Even though the vast majority of users would immediately delete the mail or not click on the phishing links, the negative experience caused residual concern about whether or not PayPal had given up their email addresses or otherwise caused them to be a target.

A much smaller group of users would give up their private financial data (including their passwords), leading to financial loss for PayPal and our customers. And of course, PayPal's financial loss was the fraudsters' profit – thus encouraging even more phishing mail and perpetuating the cycle. No matter which path the customer followed, PayPal observed consistently lower activity from these users.

1.3 Our Response

We determined five blocking points (illustrated in Diagram 2: Five Blocking Points) where PayPal could break the business model of fraudsters, and we applied a strategy against each point. There would not be one silver bullet, but perhaps five would do the trick.

Diagram 2: Five Blocking Points



Block / Strategy	Short Explanation
1. Reclaim email	Prevent phishing mail from ever entering customers' inboxes. To be judged by the adoption of email authentication by ISPs who agree not to deliver unsigned email.
2. Block phishing sites	Prevent phishing sites from being displayed to customers.
3. Authenticate users	Prevent stolen login/password combinations from being used on PayPal.com.
4. Prosecute	Create a disincentive by pursuing legal prosecution (in partnership with government and law enforcement).
5. Brand & Customer Recovery	Ensure that targeted customers would still use PayPal.

This paper concentrates on our top two strategies of reclaiming email and blocking phishing sites. For each, we will address both our "passive" and our "active" solutions, corresponding to the different segments of our user base. There are also a few other strategies that we believe are important, and while they may be ancillary to those described above, we'll give some detail about them nonetheless.

2.0 Definitions, Scope and Principles

Crimes seldom occur in isolation, and the same is true of phishing. By its strictest definition, phishing is only the act of "casting a net of impersonating email in order to steal log-in credentials." Of course, phishing is only the first step in a series of crimes that also frequently includes unauthorized account access, identity theft, financial fraud, money laundering, and others. While all of these crimes need to be considered together, this paper focuses primarily on the first step of this crime – phishing, specifically via email.

"...phishing is only the first step in a series of crimes that also frequently includes unauthorized account access, identity theft, financial fraud, money laundering, and others. While all of these crimes need to be considered together, this paper focuses primarily on the first step of this crime – phishing, specifically via email."

“If phishing never makes it into a customer’s inbox, the customer cannot become a victim.”

Principles

1. No Silver Bullet

We have not identified any one solution that will single-handedly eradicate phishing; nor do we believe one will ever exist. Instead, our approach relies on a holistic “defense in depth” model. In this approach, there are multiple layers of defense – while no single layer can defeat phishing on its own, in tandem they can make a huge difference, with each layer shaving off some percentage of crime that otherwise would have occurred.

2. Passive Versus Active Users

When thinking about account security, we observe two distinct types of users among our customers. The first and far more common is the “passive” user. Passive users expect to be kept safe online with virtually no involvement of their own. They will keep their passwords safe, but they do not look for the “s” in https nor will they install special software packages or look for digital certificates. The active user, on the other hand, is the “see it/touch it/use it” person. They want two-factor authentication, along with every other bell and whistle that PayPal or the industry can provide. Our solution would have to differentiate between and address both of these groups.

Note that for this paper, our use of the words “active” and “passive” describe customers’ views of protecting themselves on the Internet. These words do not indicate the frequency with which customers use their PayPal accounts.

3. Industry Cooperation

PayPal has been a popular “target of opportunity” for criminals due to our large account base of over 141 million accounts² and the nature of our global service. However, while we may be an attractive target, we are far from the only one. Phishing is an industry problem, and we believe that we have a responsibility to help lead the industry toward solutions that help protect consumers – and the Internet ecosystem as a whole.

4. Standards-based

A preference for solutions based on industry standards could be considered a facet of industry cooperation, but we believe it’s important enough to stand on its own. If phishing is an industry problem, then industry standard solutions will have the widest reach and the least overhead – certainly compared to proprietary solutions. For that reason, we have consistently picked industry standard solutions.

3.0 Reclaiming Email

3.1 Email Signing and Blocking

If phishing never makes it into a customer’s inbox, the customer cannot become a victim. Therefore, our number one strategy centered on a creative use of new email signing standards and cooperation with major Internet Service Providers (ISPs) to actually block unsigned email that looked to be from PayPal – before the mail reached the customers. The tough question, though, was how?

For years, the industry had spoken of creating digital signatures in email messages that could tell the end user exactly who sent the message. This was a fine solution for “active” users, but it would never work for the passive group.

As we looked at email signing, we realized that it needed to go one step further than just signing email. In order to be truly effective, ISPs had to actually throw the fraudulent email away. Essentially, our theory was:

- PayPal signs every outbound email.
- ISP scans every incoming mail to see if the from field is set to “@paypal.com”.
- If yes, then ISP checks the email signature to see if it matches our published key.
- If yes, deliver.
- If no, drop the email.

A key difference in this approach is that the email would be dropped at the ISP’s network edge rather than delivered or put in the customer’s spam folder. From PayPal’s point of view, even a spam phishing was a poor customer experience.

While we thought this solution could be effective, it would unfortunately require every ISP and every phishing-targeted company to create individual agreements – a highly unlikely situation. Instead, we concluded that our work here would be in two phases – an experimental phase, where we proved the concept, and then an implementation phase, where we tried to spur large-scale industry acceptance.

We therefore looked to find one ISP partner to prove the signing and blocking concept. If correct, we would figure out the implementation details for the industry-wide solution later. While initially daunted by the sheer number of ISPs, we quickly realized that six of the world's biggest ISPs cover more than 80% of our customer base.

In mid 2006, we began working with Yahoo! Mail to implement our solution. Yahoo! proved to be a fine match, not only because they were ten miles down the road from us, but also because Yahoo! was prominent in that list of six. In addition, Yahoo! played a significant role in the development of DomainKeys, one of the technologies we used to sign all outbound email. By mid 2007, we were using DomainKeys and SPF (Sender Policy Framework, another technology that works against spam email) and working with Yahoo! on the blocking rules.

After going live with DomainKeys email blocking in October of 2007, we have seen impressive results. In the first few months we successfully prevented the delivery of more than fifty million phishing messages from reaching the inboxes and bulk folders of unsuspecting consumers. Perhaps just as exciting is the fact that we've also seen a significant drop-off in the number of attempts to spoof PayPal in Yahoo! Mail, meaning far fewer fraudsters even try to send these scams to Yahoo! Mail users.

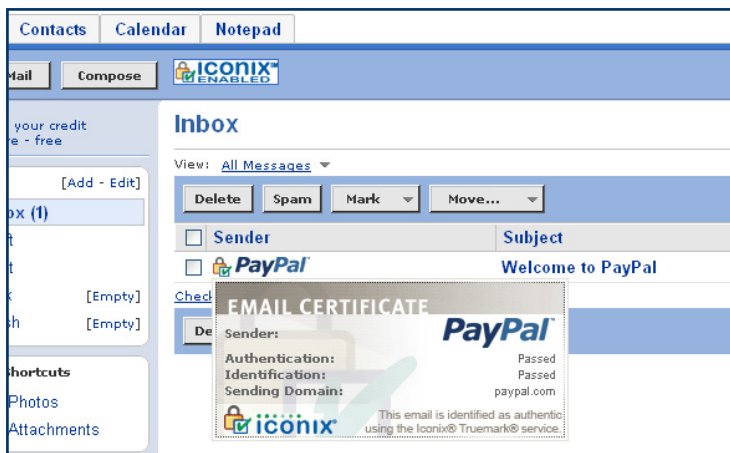
3.2 Visual Identification

Our original theory was that email signing could also help us deliver a solution for our active users – those who wanted to participate in their security – by rendering the digital signatures for astute users to recognize.

Until all ISPs enforce DomainKeys and SPF, there will be gaps in the protection that email signing and blocking cannot solve. Therefore, the second half of our email strategy is to work with the providers of email clients – of which there are also a relatively small number – to ensure that the signatures which are embedded in email are recognized by these clients.

To reach the active users who do not access their email through a signature-rendering email client, PayPal started working with Iconix, which offers the Truemark plug-in for many email clients. The software quickly and easily answers the question of “How do I know if a PayPal email is valid?” by rewriting the email inbox page to clearly show which messages have been properly signed. See diagram below.

Diagram 3: Iconix Digital Signature in Yahoo! Mail Client



“In the first few months we successfully prevented the delivery of more than fifty million phishing messages from reaching the inboxes and bulk folders of unsuspecting consumers.”

“One of the most effective things that can be used to help protect consumers are anti-fraud warning pages that warn them when they click on a link that is a confirmed phishing site.”

4.0 Block Phishing Sites

The next defensive layer assumes that: 1) phishing mail has entered a customer’s email inbox; and 2) the customer has clicked on the link. Many current browsers³ contain features that help protect consumers, and we refer to these as “safer browsers.”

4.1 Unsafe browsers

There is of course, a corollary to safer browsers – what might be called “unsafe browsers.” That is, those browsers which do not have support for blocking phishing sites or for Extended Validation Certificates (a technology we will discuss later in this section). In our view, letting users view the PayPal site on one of these browsers is equal to a car manufacturer allowing drivers to buy one of their vehicles without seatbelts.

The alarming fact is that there is a significant set of users who use very old and vulnerable browsers, such as Microsoft’s Internet Explorer 4 or even IE 3. Inevitably, this set of users is a subset of the passive group. We argue that it’s critical to not only warn users about unsafe browsers, but also to disallow older and insecure browsers. Further, we suggest that any Web site that asks for personal or financial information should consider logic along the following lines:

Version N (current) – allow with no messaging.

Version N-1 (previous major version) – allow, but with a warning message.

Version N-2, or older – disallow, with a message indicating why.

At PayPal, we are in the process of re-implementing controls which will first warn our customers when logging in to PayPal from those browsers that we consider unsafe. Later, we plan on blocking customers from accessing the site from the most unsafe – usually the oldest – browsers.

4.2 Blacklists

Blacklists form the core of the mechanism behind interstitial warning pages, which warn customers when they are about to enter personal or financial information on a site that is or may be a fraudulent Web site. Thus, to the extent that the blacklists are current, complete, and timely, they work very effectively. If they are not, then blacklist performance is less adequate.

Quite a bit of blacklist processing occurs “behind the scenes” in industry organizations such as APWG, in commercial organizations such as MarkMonitor, and in major IT companies such as Microsoft and Google.

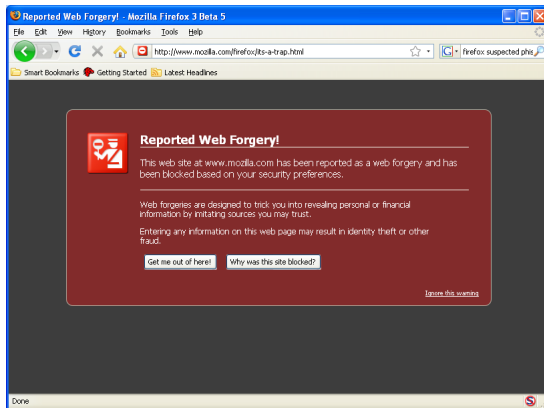
However, PayPal also generates quite a bit of blacklist data directly. Through customer education we have asked our customers to forward potentially fraudulent email to spoofoff@paypal.com. Our security team then uses an automated system to analyze the email to 1) determine whether or not it is legitimate; 2) respond appropriately to the customer; and 3) if it is fraudulent, extract all of the URLs from the email and insert them into our own blacklist system. We also supplement this list with a commercial blacklist feed service to help us optimize our spoof email handling, as well as to initiate requests to shut down the fraudulent sites.

We have spent a significant amount of time re-engineering this system so that we can confirm a site as quickly as possible (to avoid “false positives”), as well as to rapidly distribute these blacklist entries to the Internet community. Currently, over 50 organizations subscribe to our Fraudcast feed as we work together to fight fraud in near real-time. Since one of the primary goals of the criminals is to persuade victims of the urgency in responding, speed is of the essence in blacklist distribution and management. We suggest that companies think about implementing a “spoof” address, as we have done, and we’d like to reiterate the value of linking that to generation of blacklist URLs.

4.3 Anti-fraud Warning Pages

One of the most effective things that can be used to help protect consumers are anti-fraud warning pages that warn them when they click on a link that is a confirmed phishing site. Technically, this is described as an interstitial warning page and tells the user, in no uncertain terms, that the page they are attempting to access is not the site they think it is. These warning pages operate through a combination of “white lists” (of known good websites), “blacklists” (of known spoof websites), and heuristic scoring models that process the available data in HTML and JavaScript before the page renders. A sample of Firefox’s warning page can be seen in Diagram 4.

Diagram 4: *Firefox's Warning Page*



“Blocking offending sites works very well for passive users. However, we knew we needed to provide visual cues for our active users in the Web browser, much like we did with email signatures in the mail client.”

At the start of this paper, we mentioned that industry cooperation and solutions for passive users would be driving themes. They both play a strong role in our strategy to block spoof sites. Three years ago, eBay (PayPal's owner) released a downloadable tool to block eBay phishing sites called the eBay Toolbar. It had advanced blacklist/whitelist technology built into it and did a good job of warning users when they were at a site that appeared to be mimicking eBay or PayPal. In other words, the technology itself worked.

The eBay Toolbar proved to be very effective for those “active users” who wanted to participate in their own security and would go through the process of downloading and installing the software. The passive users, unsurprisingly, ignored it.

To reach the huge passive user base, we needed wider distribution. Working with Microsoft, we helped promote the concept of site blocking, which Microsoft then included in its IE 7 release. Therefore we could simply ride the curve of adoption as IE 7 proliferated across computers around the world. Just over a year after its release, IE 7 already achieved adoption rates of more than 40% of our users. This is an excellent example of getting the right technology built into the underlying platform that consumers use without requiring separate action by the consumer. Simply the act of buying a new computer provided added security for the user.

4.4 Extended Verification SSL Certificates

Blocking offending sites works very well for passive users. However, we knew we needed to provide visual cues for our active users in the Web browser, much like we did with email signatures in the mail client.

Fortunately, the safer browsers helped tremendously. Taking advantage of a new type of site certificate called ‘Extended Validation (EV) SSL Certificates,’ newer browsers such as IE 7 highlight the address bar in green when customers are on a Web site that has been determined legitimate. They also display the company name and the certificate authority name. So, by displaying the green glow and company name, these newer browsers make it much easier for users to determine whether or not they're on the site that they thought they were visiting.

PayPal was one of the first companies to adopt EV Certificates. More or less all of the pages on our site are SSL encrypted, and they all use EV Certificates. And after nine months of usage, PayPal's data suggests that there is a statistically significant change in user behavior. For example, we're seeing noticeably lower abandonment rates on signup flows for IE 7 users versus other browsers. We believe that this correlates closely to the user interface changes triggered by our use of EV certificates.

5.0 Ancillary Strategies

5.1 Customer Education

If there is a single “silver bullet” to preventing phishing and other forms of e-crime, it is indeed customer education. But as we noted right at the beginning of this section, it's also incredibly hard

“If there is a single ‘silver bullet’ to preventing phishing and other forms of e-crime, it is indeed customer education.”

to achieve, due to the sheer numbers of Internet users and varying levels of knowledge about online safety. Nonetheless, we encourage all organizations to do what they can to help educate their customers – tailored specifically to their needs.

As of early 2008, there are somewhere in the order of 1.1 billion users of the Internet, with another billion expected within the next decade. Of this 1.1 billion, the vast majority have been using the Internet for less than a decade.

Repeated studies show that the vast majority of Internet users do not know how to protect themselves online. We strongly believe that the majority of these users fall into the passive group.

The e-commerce industry therefore bears a heavy responsibility to ensure that we collectively educate Internet users about how to protect themselves and their computers. PayPal believes that this is a long-term responsibility, and it is an area in which we’ve invested heavily. At this point, there is a substantial amount of material available for companies to help educate their customers. Historically, we have found these to be most useful:

- StaySafeOnline: www.staysafeonline.org
- StaySafe.org: www.staysafe.org
- GetSafeOnline: www.getsafeonline.org
- FTC: www.ftc.gov/bcp/menus/consumer/tech/privacy.shtm
- WHO@: www.haltabuse.org/resources/online.shtml
- WindowsLive: <http://onecare.live.com/site/en-us/article/websurfsafe.htm>
- SafeKids: www.safekids.com/child_safety.htm
- US Computer Emergency Readiness Team: www.us-cert.gov/cas/tips/ST06-003.html
- McAfee: us.mcafee.com/root/identitytheft.asp?cid=23903
- Symantec: www.symantec.com/norton/cybercrime/prevention.jsp

We also view the PayPal Web site as a place to educate customers about how to stay safe online, and we’ll continue to invest heavily in this area going forward. www.paypal.com/securitycenter

Since different users have different levels of sophistication in their knowledge of the Internet, we customize our messages to various PayPal customers. For the technically aware active user, we explain how to use “hovering over” a link to expose the underlying URL. They can then look at the link and determine that a link to <https://www.paypal.com/foo> is probably safe, whereas <http://www.badguy.ru> is almost certainly not PayPal.

For the technically naïve, we have a rather different series of messages:

- 1) Don’t click on links in email.
- 2) Always open up a fresh browser window.
- 3) Go directly to the Web site of the organization concerned.

The above messages would – if universally followed – kill phishing.

As discussed earlier in this paper, we encourage our customers to forward suspicious email to us at spoofoff@paypal.com. (There’s a similar address at eBay: spoofoff@ebay.com.) These addresses are linked to a fairly complex database and rules engine, and respond quickly to customers with, “Yes, this email is legitimate and from us;” or “No, this email is fraudulent; please delete it.”

We find that this accomplishes several things:

- 1) It helps customers feel personally connected to fighting crime.
- 2) It helps them differentiate reliably between legitimate and phishing email.
- 3) It gives us a great source of known phish site URLs which can then be directly utilized in various ways.

Because of the value that this specific approach has, we strongly recommend that all organizations consider implementing a similar spoofoff@yourcompany.com reporting system.

5.2 Site Shutdown

As with many of the other defensive layers that we describe in this paper, it’s intuitively obvious that, if a customer attempts to browse to a spoof site, but that site is unavailable, they will not be victimized. It goes without saying that the criminals are trying to engender a (false) sense of urgency in their communications, so time is of the essence in executing a site shutdown.

This is a tricky topic. In most of the West, getting a site shut down can be accomplished within a few hours. However, this process has some weak spots and can be improved with judicious attention to detail. We described earlier the process by which candidate URLs need to be manually reviewed for inclusion within industry blacklists. This is a good starting point for process re-

engineering to ensure that the reviews happen as quickly as possible, so that a particular URL can be confirmed as a spoof site.

First, once a site has been determined (by manual review) to be a malicious spoof site, it's important to request site shutdown through the hosting ISP as soon as possible thereafter. The vast majority of ISPs have a well defined site shutdown request process, and it is fairly simple to submit those requests. The trick is that one needs a system to keep track of whether a particular site is still available, and to escalate to the responsible ISP in a more manual fashion, typically by simply calling them.

Unfortunately, some ISPs have very limited off-hours support and so in some countries, ISPs may not respond at all to shutdown requests over the weekend, for example.

Second, in some cases, if the ISP is unresponsive, it is sometimes possible to get the domain shut down via the domain registrar. This is highly variable. Some registrars are extremely responsive and, if provided with appropriate documentation, will remove the DNS entries for a spoof domain quickly. Other registrars are only minimally helpful, and will react only after days or weeks – far too late to make any difference in preventing consumers from being victimized.

5.3 Authentication

One of the reasons that criminals have been relatively successful in their phishing campaigns is that most sites still have weak authentication – that is, a user ID/password is sufficient to gain access to site services. Worse, account recovery flows are typically designed to use so-called “secret” information to validate customers. Unfortunately, many phishing attacks also garner that data too, and so it is now relatively easy for many criminals not only to impersonate a customer by their login credentials, but also via their other confidential data.

For example, many customers share passwords across multiple e-commerce websites. If a criminal either compromises one of these sites (assuming that the site doesn't follow best practices in hashing the password), or manages to successfully phish a customer to get them to reveal a password, then the criminal will often try the same password at other sites.

The only solution to this is to strengthen authentication practices, especially for access to financial accounts. Globally, this is starting to occur, but best practices for strong authentication have not yet fully emerged. Nor has this problem been solved in a way that doesn't make consumers negotiate a maze of different authentication methodologies.

In the case of PayPal, we have launched the PayPal Security Key in three markets globally – the US, Australia and Germany. In its current incarnation, the Security Key is an OATH-based token which operates on the VeriSign VIP network. We plan to launch the PayPal Security Key in other markets globally, and we're also considering additional form factors based on customer demand.

5.4 Fraud Models

In general, PayPal doesn't say a great deal publicly about our fraud-modeling technology, and we take this approach to protect the effectiveness of our techniques. We do make considerable investments in both fraud detection engines and in well-trained fraud analysts to make decisions in cases where automation cannot beat humans. PayPal's anti-fraud technology is a competitive differentiator, bolstered even further by our recent acquisition of Fraud Sciences.

As criminals innovate in their attacks against us and our customers, our fraud detection techniques have to keep up. It is an arms race in its truest form, and continual investment is required to stay ahead. Any company on the Internet needs to have an anti-fraud approach, and it can never be viewed as static.

5.5 Law Enforcement

The first thing to note is that it's important to build relationships with law enforcement agencies before they're needed. By and large, law enforcement professionals are extremely keen to engage with private industry, and so any kind of outreach is generally welcomed.

PayPal has worked hard to build relationships with law enforcement around the world, and on occasion, our agents actually help train law enforcement officials, prosecutors, and judicial staff on investigating and prosecuting e-crime cases. This approach has helped our ability to successfully find, catch, and convict e-criminals in jurisdictions across the globe.

“Any company on the Internet needs to have an anti-fraud approach, and it can never be viewed as static.”

“...our Government Relations team participates in public-private partnerships that are specifically focused on identifying and prosecuting e-crimes.”

5.6 Government Relations

A significant part of our overall program is concerned with outreach to regulators and policy makers. There are a couple of main pillars to our work here.

First, we work with regulators to educate them about the contours of e-crimes. Our Government Relations team accomplishes this by regularly meeting with global policy makers and participating in industry fora.

Second, we try to encourage the adoption of reasonable laws against e-crime. For example, in the crime of phishing, where does the crime occur?

- 1) When mailing out a large email campaign pretending to be another organization?
- 2) When operating a website which captures credentials from consumers who think that the website is run by a legitimate organization?
- 3) When logging on with those credentials to investigate whether the online balance or funding sources justifies draining funds from an account?
- 4) When actually moving funds out of the account?

In all too many jurisdictions, the crime occurs at stage four. In a very few jurisdictions, basically the US and parts of Western Europe, the crime occurs at stage one. We believe that in every jurisdiction, in any country around the world, prosecution needs to happen at stage one.

Third, our Government Relations team participates in public-private partnerships that are specifically focused on identifying and prosecuting e-crimes. Through these partnerships, PayPal works with other financial institutions and law enforcement to gather, analyze, and act on information about financial crimes.

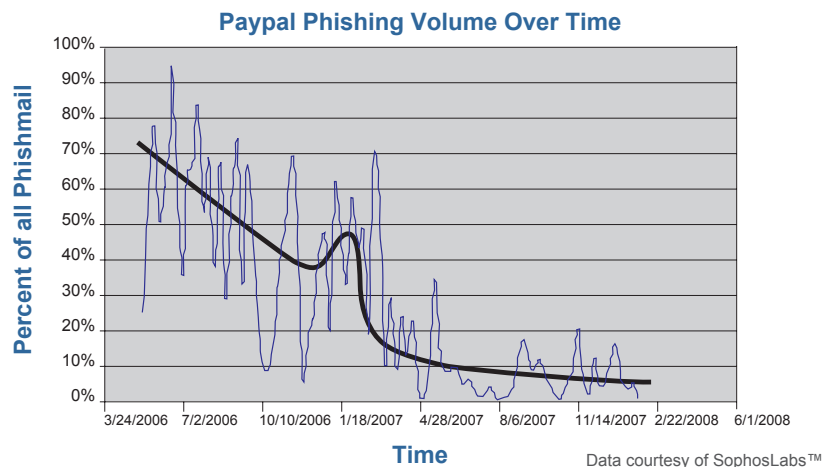
Fourth, we work with regulators to encourage them to invest more in investigating and prosecuting e-crime. A recent Gartner report⁴ claimed that the global haul from phishing is now estimated to be in the range of \$3.2 billion. This sum is almost half the size of the illegal drugs business between the US and Mexico⁵, and yet the level of law enforcement staff assigned to e-crime is trivially low.

Finally, we are starting to educate lawmakers about a very significant political dimension to the problem. That is, there are jurisdictions where the chances of being arrested in connection with e-crimes is effectively nil. In those jurisdictions, the local law enforcement authorities have a tendency to view phishing as a crime which occurs against decadent Westerners and not against their own citizens. Worse, the criminals are almost viewed as social contributors because they bring large amounts of cash into their countries. The only way that this problem will be solved is through substantial amounts of political discussion at a very senior level.

6.0 The Results

Within a few months of implementing the broad-based strategy described above, we started seeing results. The diagram shows graphically how much the phishing rate has dropped – from rates that were peaking at over 80%, to rates that are now lower than 10%.

Diagram 5: PayPal as Percent of Industry Phishmail



⁴ “Phishing Attacks Escalate, Morph and Cause Considerable Damage,” December 2007

⁵ “National Drug Threat Assessment 2006”, National Drug Intelligence Center, January 2006

These results surprised us. By this point in our test we were still operating with just one email blocking provider. Yet the results are clear. While it is impossible for us to draw a 100% causation conclusion from our strategy, we believe that the hurdles that we are placing in front of the fraudsters with respect to phishing are driving them to preferentially target other brands. Of course, that is a bittersweet result. As members of the Internet community, we would prefer that phishing was eradicated entirely from consumers' inboxes. Rather, we hope this first step that we have demonstrated is a clear call to the industry that together we can make significant progress against this crime.

7.0 Conclusions and Next Steps

While our initial test of email signing/blocking was highly successful, we realize there are great distances still to travel to make our email signing and blocking vision a widespread reality. We mentioned earlier that the approach we've used for initial deployment of signing/blocking (working with large ISPs individually) is non-scalable for the industry as a whole.

Essentially, the standards need to support mechanisms by which senders can specify policies to describe how their email should be processed by "verifiers." There is a proposal within the Internet Engineering Task Force for "Sender Signing Practices" (SSP), which would be additive to the existing DNS mechanisms. However, at this point this work is at an early stage. In practice, we suspect that a robust policy language, such as SSP, will turn out to be the best way to proceed. We strongly encourage the rapid development and formal adoption of this approach.

At the beginning of this paper we observed how our strategy to prevent financial loss was successful, but with the unintended consequence of driving more phishing. We refer to this phenomenon as 'squeezing the balloon' – an analogy to there being a finite number of fraudsters who will simply migrate within the space provided to commit their crimes.

As we look forward, with twelve months of history to build upon, we believe we can already see the 'balloon squeeze' effect at work. We have observed a rise in the number of viruses and Trojans that are deployed to steal user credentials – a far more sinister way to accomplish the same end-result as phishing. At PayPal, we are actively watching this new threat and creating strategies, tools, and techniques to fight it.

Without question, we are in a fast-moving chess match with the criminal community. Each of our moves is countered by one of theirs. We continue to believe that with vigilant efforts, and the cooperation of the entire Internet community, we can pop this balloon for the good of our companies, our customers, and the future of the Internet.

There's clearly no "silver bullet" which will deal with phishing. Rather, we've made a credible case that a multi-layered strategy, such as the one we've laid out, can in fact make a significant difference in dealing with the crime. We encourage the rest of the industry to evaluate their anti-fraud efforts and adopt a fraud prevention strategy along these lines. As the old adage goes, "united we stand; divided we fall."

References

"Over 75% of all phishing emails target PayPal and eBay users," Sophos, July 27, 2006

"Phishing Attacks Escalate, Morph and Cause Considerable Damage," Gartner Inc., by Avivah Litan, December 13, 2007

"National Drug Threat Assessment 2006," National Drug Intelligence Center, January 2006

About PayPal

PayPal is the safer, easier way to pay and get paid online. The service allows anyone to pay without sharing financial information and gives consumers the flexibility to pay in any way they prefer, including through credit cards, bank accounts or account balances. With more than 57 million active accounts in 190 markets and 17 currencies around the world, PayPal enables global ecommerce. PayPal is an eBay company. More information about the company can be found at <https://www.paypal.com>.

"We encourage the rest of the industry to evaluate their anti-fraud efforts and adopt a fraud prevention strategy along these lines."