

Improving Phishing Countermeasures: An Analysis of Expert Interviews

Steve Sheng,¹ Ponnurangam Kumaraguru,² Alessandro Acquisti,³ Lorrie Cranor,^{1,2} and Jason Hong²

¹Department of Engineering and Public Policy, Carnegie Mellon University, Pittsburgh PA 15213 USA

²School of Computer Science, Carnegie Mellon University, Pittsburgh PA 15213 USA

³Heinz School of Public Policy, Carnegie Mellon University, Pittsburgh PA 15213 USA

In this paper, we present data from 31 semi-structured interviews with anti-phishing experts from academia, law enforcement, and industry. Our analysis led to eight key findings and 18 recommendations to improve phishing countermeasures. Our findings describe the evolving phishing threat, stakeholder incentives to devote resources to anti-phishing efforts, what stakeholders should do to most effectively address the problem, and the role of education and law enforcement.

Index Terms—anti-phishing, expert interview, electronic crime.

I. INTRODUCTION

PHISHING is a widespread problem that is impacting both business and consumers. In November 2007, MessageLabs estimated that 0.41% of the 3.3 billion emails going through their system each day were phishing emails [1]. Microsoft Research recently estimated that 0.4% of email recipients are victimized by phishing attacks [2]. The annual cost to consumers and businesses due to phishing in the US alone is estimated to be between \$350 million and \$2 billion [3].

To reduce the damage due to phishing, stakeholders have implemented their own countermeasures: major web browsers have built-in filters (e.g. [5], [6], [7]), Internet service providers filter suspicious phishing emails, law enforcement officers find and prosecute phishers, and US government agencies and corporations now educate consumers on phishing.

Phishing scams have also evolved, sometimes at a faster pace than countermeasures. Phishers launch attacks on specific groups (e.g. users of social networking sites) through multiple channels (e.g. phone, instant messaging), and phishing toolkits and compromised credentials are readily available for sale at low prices on Internet black markets [8]. Sophisticated phishing schemes such as man-in-the-middle attacks and malware are becoming more frequent [9].

As the battle against phishing continues, many questions remain about where stakeholders should place their efforts to achieve effective prevention, speedy detection, and fast action. Do stakeholders have sufficient incentives to act? What should be the top priorities for the anti-phishing community? To provide insights into these questions we conducted 31 in-depth interviews with anti-phishing experts between May 2008 and May 2009. We selected experts from academia, Computer Emergency Response Team (CERT) centers, the Anti-Phishing Working Group (APWG) officers, law enforcement, and key industry stakeholders. We sought their expertise on the current and future state of phishing attacks, countermeasures that should be implemented to fight phishing

more effectively, and incentives that various stakeholders have in their fight against phishing.

The experts we interviewed agreed that phishing is evolving into a more organized effort. It is becoming part of a larger crime eco-system, where it is increasingly blended with malware and used as a gateway for other attacks. Some of the experts suggested that incentives for fighting phishing may be misaligned, in the sense that the stakeholders who are in a position to have the largest impact do not have much incentive to devote resources to anti-phishing efforts. In terms of countermeasures, experts identified improving law enforcement and shutting down money trails as top priorities. They also identified operating systems vendors, web application providers, browsers, and Internet service providers as stakeholders with key technology influence on phishing. Finally, experts agreed that education is an important factor that is not emphasized enough; however, they did not agree on the extent of the impact that education may have. We present these findings and a set of recommendations to improve countermeasures.

Although previous reports have studied phishing and issued recommendations, to the best of our knowledge this is the first study that synthesizes the opinions of experts from different fields, and examines the incentives of various stakeholders to contribute to anti-phishing efforts.

II. RELATED WORK

In response to the growing phishing problem, government agencies, industry groups, and consumer groups have conducted studies and issued recommendations [10,11,12,13].

The Financial Services Technology Consortium's report is the first report that analyzed how phishing works by articulating the life cycle of phishing. It also encouraged financial institutions to assess the costs and risks associated with phishing, develop better intelligence on phishers through improved sharing, and invest and adopt in better mutual authentication. However, the report did not issue recommendations for non-financial institutions who also have high stakes in the phishing problem [10].

Manuscript received June 30, 2009. Corresponding author: Steve Sheng (e-mail: shengx@cmu.edu).

978-1-4244-4626-1/09/\$25.00 ©2009 IEEE

Table 1: Phishing stakeholders. Primary victims suffer direct losses from phishing. Infrastructure providers have technical capabilities to mitigate the problem. For-Profit protectors sell solutions to primary victims and infrastructure providers. Public protectors include law enforcement officials, computer emergency response teams, and academic researchers.

Categories	Examples of key stakeholders	Roles
Consumers	--	Primary victims
Organizations	Military, Universities, Corporations	
Financial Institutions	Bank of America, Citibank, Paypal	
Merchants	Online merchants (eBay, Amazon), offline merchants	
Registrars and Registries	GoDaddy, Verisign	Infrastructure providers
Internet Service Providers	AT&T, Comcast, AOL, Universities	
Email Providers	Gmail, Yahoo!Mail, Hotmail	
Browsers	Internet Explorer, Firefox, Safari	
Software Vendors	Symantec, RSA, MarkMonitor, Cyveillance	For-profit protectors
Law Enforcement	Federal Bureau of Investigation (FBI), Secret Service state and local enforcement	Public Protectors
Computer Emergency Response Teams	CERT-CC, CSIRTs	
Academia		

The Identity Theft Technology Council report also analyzed different phases of phishing and recommended a set of 21 technical countermeasures [11]. We selected a subset of recommendations from this report as a starting point for discussion in our expert interviews. However, we updated the set to address non-technical countermeasures as well as new and evolving threats that were not discussed in the report. In addition to discussing the set of recommendations, we also studied the incentives that stakeholders have to implement them as well as how the incentives can be increased.

In addition to these reports, the Anti-Phishing Working Group (APWG) has issued a set of best practices and recommendations for hacked website owners [14], registrars [15], and ISPs and mailbox service providers [16]. Each of these reports focus narrowly on one particular area. In our analysis, we analyzed the phishing issue holistically and asked our experts to prioritize their recommendations based on their importance and effectiveness.

III. STAKEHOLDERS

Phishing involves many stakeholders, including consumers, financial institutions, online merchants, Internet Service Providers (ISPs), mail client and web browser vendors, and law enforcement. In this paper, we have classified stakeholders into the following categories: *primary victims*, *infrastructure providers*, *for-profit protectors*, and *public protectors*. Table 1 describes these stakeholders and their roles. We used it to select experts and structure our interviews.

Primary victims: In most cases, consumers, organizations, financial institutions, and merchants are direct targets of phishing attacks. Each of them is negatively affected by phishing in a different way. For example, consumers who fall for phishing can potentially become victims of identity theft: they not only suffer monetary loss, but also psychological costs (e.g. fear, anxiety). Organizations such as the military and corporations worry that phishing may lead to further

compromise of credentials that can be used to steal key intellectual property or conduct corporate espionage. Financial institutions lose money from fraud conducted with credentials acquired through phishing, and merchants lose money because these financial institutions eventually charge them for the fraudulent transactions. In general, these entities are most impacted by phishing, and have the strongest incentive to protect against phishing. However, as shown later in the result section, some of them have limited capabilities to counter phishing attacks.

Infrastructure providers: Internet service providers, email providers, browsers, domain name registrars, and registries are infrastructure providers. In most cases, phishers do not go after these providers for their money; instead, they seek to gain access to the entities' infrastructures so that phishers may launch their attacks. For example, phishers register fake domain names with registrars. Phishers use compromised machines from Internet Service Providers as part of a botnet to launch phishing campaigns, sending emails to end user mailboxes or compromising mail provider accounts to send phishing emails. These stakeholders are important to study, as they are in a better position than most victims to protect against phishing. However some infrastructure providers do not lose money from phishing, so they may not have sufficient incentives to devote resources to combating phishing. In our interview study, we asked experts what these stakeholders can do and examined whether or not they have incentives to do so.

For-profit protectors: Certain organizations actually benefit from phishing because it is an opportunity to develop and sell products to other stakeholders. These include companies that sell spam filters and anti-virus software, as well as companies that take down phishing websites. As they are the front-line defenders against phishing, we selected a few of our experts from these companies. However, as they make money from combating phishing, it could somewhat

bias their recommendations. We discuss these potential biases in detail in the methodology section.

Public protectors: In contrast to anti-virus vendors and spam filter companies who are for-profit protectors, law enforcement, computer emergency response teams (CERT), and academics are public protectors.

There are some para-organizations such as the Anti-Phishing Working Group (APWG) and the Message Anti-Abuse Working Group (MAAWG) that aim to bring different stakeholders together to fight more effectively against phishing. Some of the experts we interviewed hold positions in these organizations. However, we did not consider these organizations as separate stakeholders in our analysis.

IV. METHODOLOGY

During May 2008 and May 2009, we conducted in-depth interviews with 31 experts involved in phishing countermeasures. In this section, we discuss how we selected the experts, the interview process, and the steps taken to analyze the data.

A. Recruitment and Participants

We recruited experts in several ways. First, we compiled a list of frequent speakers from 2004 through 2007 APWG member conferences and generated a list of well-known experts in academia and industry. To recruit law enforcement officers, we attended the 2008 Digital PhishNet conference. To recruit experts in Internet service providers, registrars, and technology vendors, we solicited recommendations from APWG's Internet Policy Committee (IPC), which is composed of 90 members from various stakeholders. Finally, we recruited additional interviewees through our own network of contacts. In order to obtain a variety of views, we tried to select candidates from different organizations who worked at different levels of company hierarchy.

We recruited a total of 31 experts responsible for, or knowledgeable of, operational or policy decisions with regard to phishing and malware prevention in their organizations. Most of the interviewees chose to remain anonymous. Table 2 shows the organizational profiles of these experts. 67% of the experts interviewed had a technical background, 20% had a policy or business background, and the remainder had a background in law or law enforcement.

In addition to the 31 experts interviewed, we also had a short interview with a legal expert on the question of liability for false positives.

The sample size of 31 balances the resource-intensive demands of in-depth interviews and analysis against the marginal return of new insights from additional participants. We had multiple participants who shared similar views on most of the topics we discussed in our interviews, suggesting that theoretical saturation was likely achieved, even with our small sample.

B. Interview Protocol

We used a semi-structured interview protocol. The protocol allowed us to ask structured questions that enabled comparable responses across participants, while providing the interviewer flexibility in drilling down on areas of particular relevance to each participant [17].

Table 2: Anti-phishing experts interviewed. For confidentiality purposes, all participants are anonymized.

Affiliation	No. of Experts
CERT	4
Academic researchers	5
APWG officers	3
Law enforcement	5
Registrars, Registries	3
Financial institutions	4
Internet service providers	3
Browser vendors	1
Other experts	3
Total	31

Each interview typically lasted 60 minutes (min =25, max = 90) and was recorded for transcription. Some interviews were conducted in-person, while others were conducted over the phone. We began each interview by asking each expert to describe his or her background and responsibilities. We then asked a set of open-ended questions about how phishing impacts their organizations, amount of losses, current and future state of phishing, and the effectiveness of current countermeasures. We then asked them specifically to comment on a set of 31 recommendations broken into six categories that we compiled through our research. Experts prioritized the recommendations in each category and provided feedback on them. Finally, at the end of each interview, we asked experts to provide additional recommendations, and if they did, we summarized and added them to our list of recommendations and asked experts about them in subsequent interviews.

C. Analysis

After completing each interview, we transcribed the audio recordings and recorded the answers to each question in a spreadsheet. We then analyzed the interview results and synthesized a series of findings and accompanying recommendations.

In our analysis, we synthesized experts' opinions by selecting themes that recurred most frequently across all interviews. We also report some of the comments that were discussed by only one or two experts, but that we found particularly useful in thinking about phishing countermeasures.

D. Limitations

Before turning to the empirical findings, it is important to note the scope and limitations of this study.

Most of the experts interviewed were from the US, but we also had some from Japan, Hong Kong, Italy and Britain. Thus, while there is some international representation, for the most part these interviews represent a US-centric view.

It is also reasonable to assume that this set of interviewees was influenced by some degree of self-selection. Registries, for example, are more likely to respond favorably to an

Table 3: High-level findings.

Categories	Findings
Evolving threat	A. Phishing is evolving to be more organized and targeted. It is becoming part of a large crime eco-system. B. Phishing and malware are increasingly blended together.
Stakeholder incentives	A. Stakeholders have varying incentives to fight phishing. B. Sometimes stakeholder incentives are misaligned.
What stakeholders should do	A. Operating systems vendors, web application providers, browser vendors, and Internet service providers are stakeholders with key technology influence over phishing. B. Organizations are conservative about filtering and warning about phish because they are worried about false positives. C. Registries and registrars can play an important role in fighting against phishing.
Law enforcement and education	A. Law enforcement should be emphasized; but law enforcement lacks the necessary tools, personnel, and resources to catch phishers. B. Shutting down money trails is very important to defeat phishers.

interview request about their phishing countermeasures if they have policies in place that are at least on par with other registries, if not better. With that said, some of the organizations we interviewed are *not* known for having outstanding records with regard to phishing.

Our findings reflect how stakeholders themselves describe what they are doing and why. In other words, we report on the perceptions of the interviewees, not the independent assessment of their actions and the factors driving them. Whenever possible, we did crosscheck information provided to us against the information from other interviews and against publicly available data, such as reports, surveys and research publications.

In addition, the interviewees are not experts in all areas, and they have biases of their own. For example, take-down vendors are more likely than others to recommend that more efforts should be focused on take-downs. We address this in a few ways. During our interviews, we let interviewees select the two to three areas in which they are most experienced to comment on. Whenever possible, we asked them to provide evidence to support their positions and recommendations, and in some instances, we tried to probe experts further by presenting a counter-argument for experts to respond to.

Despite these limitations, our approach is an important complement to purely technical analysis of phishing (e.g. [11]). First, our interview approach synthesizes the opinions of experts from many different fields. It would be difficult to obtain this information through other methods. Second, our interviews examine the incentives of various stakeholders to

contribute to anti-phishing efforts, an important consideration in producing workable solutions. For example, past qualitative research in information security investments has proven to be a valuable complement to the knowledge generated through quantitative modeling or analysis (e.g. [18],[19]).

In the next sections we present the findings from our interviews. We classified our findings into four topical categories: the evolving threat, stakeholder incentives, what stakeholders should do, and law enforcement and education. We also provide a set of recommendations based on these findings. Table 3 presents the high-level findings from the interviews.

Finally, this paper does not discuss some relevant technologies such as email authentication (SPF, DKIM), extended validation certificates. We found no consensus on the effectiveness of these technologies.

V. EVOLVING THREAT

A. Phishing is evolving to be more organized and targeted. It is increasingly used as a gateway to other attacks.

We asked experts to describe the phishing attack trends they have observed and predict how phishing attacks will continue to evolve. Experts observed that phishing attacks are becoming more organized. One technical expert in law enforcement explained:

These are criminal organizations that exist that perpetrate these types of fraud. It is not likely your teenage hacker like in the old days. They are criminal organizations with business plans and contingency plans. They are typically involved in other crimes besides phishing. It could be malware, it could be hosting other content, possibly child pornography, and it could be the old 419 scams and mule schemes. What we see is that these types of folks don't just do one thing. They either do other things or work with groups that do other things.

One example of an organized group mentioned frequently by experts is the rock phish group, which is believed by many experts to originate from a group of phishers in Eastern Europe. One academic researcher said 88% of the one million URLs his research group processed in October 2008 had rock phish characteristics. Published studies have also analyzed the frequency of fast flux phishing attacks. For example, Moore et al. found that 68% of the phishing emails in their study sample were sent using fast flux techniques [20].

Another trend that experts observed is that phishing is increasingly used as a gateway to other attacks. One expert from a major browser vendor said:

We are seeing a lot of blended attacks, where a piece of the infrastructure is a phishing attack, but that's not necessarily the end goal.... It is malware, it's affiliate advertising, it's spam as form of advertising, scams, and ring tones, there is a number of ways to monetize. But the goal is to look for not only the traditional stuff but ways to monetize groups of users. And you know,

stealing a password is a pretty good way to tag into real people, real networks, so we see the social network site is being targeted very heavily, and it's the result of that.

One of the experts from a major US bank agreed, and added that his institution had been seeing an increasing amount of cross channel fraud, where credentials harvested through traditional phishing attacks were being used to commit fraud in other channels such as telephone banking.

Finally, experts agreed that phishing attacks are evolving into be more targeted attacks, which are very effective and harder for spam filters to detect. Recent phishing attempts to defraud top executives are examples of these targeted attacks. Past research has demonstrated the effectiveness of spear phishing attacks. For example in a study at Indiana University, 16% of participants fell for regular phishing emails, but 72% fell for spear-phishing emails [21].

Phishers kept moving to new targets as traditional targets of phishing attacks have devised response plans. Some experts thought that small and medium brands would become the next victims. Others speculated that credit unions, social network sites, and Xbox live accounts would be increasingly targeted.

B. Phishing and malware are increasingly blended together.

Experts mentioned that malware attacks that use phishing emails are on the rise and pose a serious threat. One academic researcher framed phishing and malware as different expressions of the same problem. He said:

You will see social engineering aspects of malware and high automation aspects of phishing. At some point, it might be hard to tell them apart ... To the attackers, it doesn't matter what they use. They know social engineering has an effect on the end user, they know script and code and have some effect on the user's machine. It is just a matter of putting what they know and what they have.

Some of the experts we interviewed believe that malware now poses a bigger threat than phishing. Their reasoning is that due to vulnerabilities in operating systems and web applications it is easy for computers to get infected with malware, and that even security-conscious users may have difficulty avoiding infection.

VI. STAKEHOLDER INCENTIVES

A. Stakeholders have varying incentives to fight phishing.

We asked experts how phishing impacts their organizations. Their responses provided insights into their organizations' incentives to fight phishing.

In general, we found that the primary victims have incentives to invest resources to protect against phishing as they suffer direct losses from phishing. Nonetheless, there is evidence that not all potential primary victims have made this investment. One expert from academia said that many midsize and smaller banks he talked to did not have a strategy for

phishing, as they had never been targets: "There is low chance that those banks are targeted, but if they are targeted, they could lose a lot of money."

The stakeholders who do invest in anti-phishing protection sometimes feel that they are carrying a disproportionate share of the burden. One expert said:

After speaking to many service providers such as financial institutions, there is one thing that stands out very clearly, a sense of "injustice," that they are often carrying the cost for something they have no ability control or even measure. For example, financial service providers, they are not able to determine if their clients, the end users, have appropriate anti-virus software or not. So one way to align the incentives is for service providers be able to audit the security posture of user clients.

Our interviews revealed information on the incentives of several types of stakeholders, described below.

Financial institutions. Financial institutions are among the primary victims of phishing as they lose money from fraud committed with compromised accounts. Currently, over 79% of phishing attacks target financial institutions [22]. A major US bank told us that over the past 12 months, their loss due to phishing and malware was \$4.5 million, accounting for 25% of their fraud loss through online channels.

Financial loss and regulatory oversight are both drivers for adopting anti-phishing technologies. One electronic fraud risk manager from a major bank in Asia mentioned that their loss to phishing and electronic crime is less than 1% of their overall fraud loss. However, they still invest a lot of money in anti-phishing efforts because regional regulators demand two-factor authentication and require comprehensive analysis for electronic crime incidents. Thus, stakeholder incentives may vary depending on local regulations.

Finally, reputation was also mentioned by some as a factor. This same risk manager mentioned that another major reason his bank was spending a lot of money in this area was that bank management wanted to position their electronic banking service as the safest in the region.

It is worth noting the inherent difficulty of obtaining accurate phishing loss figures for financial institutions. It is difficult to separate phishing from other electronic fraud, such as malware. Furthermore, such losses impact a variety of different parts of a company, such as customer service, and thus may not be fully accounted for by the fraud department. Finally, it is difficult to quantify indirect loss such as damage to one's reputation.

Even if financial institutions have accurate phishing loss estimates, they often do not have incentives or regulatory requirements to disclose them. They may prefer not to disclose these losses due to fear of brand erosion due to negative publicity. This leads to a wide range of loss estimates that differ by an order of magnitude (e.g. [3] vs. [4]).

Merchants. Merchants lose money because financial institutions eventually charge them back for fraudulent transactions. When a phisher makes a purchase using a stolen

credit card, the credit card company usually charges the merchant for the loss. With online and telephone transactions known as “card-not-present” transactions, merchants assume this liability directly if cardholders dispute a charge. The Merchant Risk Council estimates that merchants who manage their risk well still lose about 1% of their revenue to credit card fraud [23].

Internet Service Providers: The ISPs we interviewed all considered phishing as part of the spam problem, which is their number one concern. Since phishing usually represents less than 1% of the spam they receive, their typical response is to filter out phish with spam. For example, one University ISP expert said, “We filter as much as we could and we would like [our users] not be sending their credit card and social security numbers online, but we don’t see that as our responsibilities to protect those numbers, it is their personal data to protect.” Other experts from academia echoed this sentiment as well.

ISPs do have an incentive when phishing targets their own mail systems. These phishing attacks typically seek to compromise users’ webmail accounts hosted by these ISPs and use them to send out more spams. ISPs have the incentive to ensure mail flows properly and avoid having their mail servers being blocked by blacklists.

When it comes to fixing compromised machines that are often used as part of a botnet to send out phishing emails, ISPs currently do little. These compromised machines sometimes form a fast flux network, in which a domain name that phishers use has multiple IP (Internet Protocol) addresses assigned to it. The phishers switch those domains quickly between the addresses (often compromised machines) so that it is not as easy to find or shut down the phishing sites. One expert from a major US ISP recognized that compromised PCs cause major problems, and told us that close to 10% of their customers’ machines were infected with malware. However, when asked why his company does not remove these computers from the network he said, “Well, they are paying [a monthly fee] ... for Internet access.”

Experts from other ISPs made similar comments and noted that fixing infected computers can be costly. Infected computers may need to have their operating systems reinstalled. One expert from an ISP mentioned that customer service is the largest cost for the ISP. However, most experts who did not work for ISPs identified infected machines on ISP networks as a major problem that needs to be fixed.

Domain Registrars: Registrars have been generally regarded as lagging in terms of phishing countermeasures. One expert claimed that registrars actually have a disincentive to fight phishing as criminals still pay them for registering phishing domains. However, another expert familiar with the registrars disagreed, saying, “Registrars would get charge back eventually because phishers are usually using fake credit cards to register these domains.” Some other experts suggested that registrars lacked the capability to detect and shutdown phishing fraud, as they work on small profit margins.

B. Stakeholder Capabilities and Incentives are Often Misaligned.

Economists have suggested that liability should be assigned

to the party that can do the best job of managing risk [24]. However, throughout the interviews, we found that the party that can do the best job is not always managing the risk.

For example, in Asia, if banks can prove that a customer acted with negligence, the bank is not liable for a phishing loss. The difficulty is to prove that customers acted with negligence. One participant from a major bank in Asia said that when his bank was first attacked by phishers, the bank reimbursed victims. However, he said, “We’ve then since done a lot of education and we have joined the association of banks for a series of community bank education programs. After that, if customers do not pay attention to the education, we consider that to be negligent, so we try not to reimburse them. Of course, if the customer starts to yell and complain to the regulators, then it is entered into a fueled debate.”

As another example, experts mentioned that merchants are held liable when phishers use fake credit card credentials to buy goods from them. When banks find out about the fraudulent charges, they will charge the merchant for it and sometimes also charge fines. This liability can be shifted if merchants implement the “Verified by Visa” program, but many merchants do not because of usability concerns. Furthermore, one expert argued that it is very difficult for merchants to notice that a credit card is stolen, noting that banks are at a much better position to make that judgment because they possess more information about the given credit card and a history of the transactions that make it easier for them to spot fraudulent charges.

As a third example, some experts claimed that ISPs are in the best position to protect their network and clean up compromised machines, but are not willing to take proactive measures because they would incur high costs while generating little benefit. One expert said:

The ISP is in a good position to inspect and identify some machines that are sending out spam and launching denial of service attacks.... There are quarantine devices that exist.... ISPs have it, but even for the ISPs using them, it is not used much. It is expensive for ISPs. If you put the user on quarantine, you end up having high customer cost, the person will call the help desk, and you have to walk them through everything. The benefit to the ISP is very low compared to the cost. This is because the ISP did not bear the cost of compromised machines, putting externalities, hosting spam, it is not infecting the ISPs bottom line, but it is impacting every one else’s bottom line.

We asked experts to comment and prioritize on a set of recommendations on the issues of incentives. We discuss the first recommendation with our experts and introduced the second recommendation based on our findings.

Recommendation (R1): Financial institutions should produce more accurate estimates of phishing losses and report these statistics.

As we mentioned earlier, accurate estimates of the phishing threat are difficult to come by, but very important. For example, it is difficult for law enforcement to open cases if

they do not have a good idea of the amount of loss or the type of damages. Similarly, without quantifying damages, it is hard for corporations to manage the risks.

For a corporation to obtain these figures, experts suggest two possible steps: first, law enforcement should collect and preserve forensics data when the phishing servers or drop accounts (email accounts used to gather stolen credentials) are seized, provide detailed information about the accounts stolen, and collaborate with banks to double check these fraud cases. Second, fraud managers within the organization should look at the organization as a whole when estimating damages, not just the online fraud itself. For example, they could examine how phishing increases customer service costs.

Financial institutions also need to report overall statistics on phishing. Currently, many financial institutions do not have incentives to report estimates of phishing losses, and fear of negative publicity serves as a disincentive. One way to address this is anonymous reporting, such as in the case of the UK payment association (APACS), which encourages its members to report their losses and aggregate them together. Another way to address the problem is that the relevant US regulator institute some form of mandatory reporting policy for fraud statistics.

Recommendation (R2): Regulators and academic researchers need to investigate the issue of incentives further.

As mentioned in our findings, some stakeholders (such as consumers or merchants) are not really equipped to protect themselves against fraud, so placing the liability or burden of proof on them would do little to help fight against phishing. On the other hand, ISPs who are in a better position to clean the compromised machines do not have incentives to do so. Further research is needed to develop incentive models and determine where incentives are misaligned and ways to realign the incentives.

VII. WHAT STAKEHOLDERS SHOULD DO

A. Experts identified operating system vendors, web application providers, browser vendors and Internet service providers as stakeholders with key technology influence over phishing.

Experts identified operating system vendors, web application providers, browser vendors, and Internet service providers as being best positioned to fight phishing.

Operating systems are crucial because their security or insecurity has far reaching effects. Experts generally praised Microsoft for their efforts in hardening their operating systems, but pointed out more to be done in this area. They gave a few recommendations that we will cover in the later part of this section.

Experts pointed out the insecurity of web applications as a significant hurdle. One technical expert charged web application vendors for the current state of the problem:

[Phishers] are losing on the email; the majority of the places are running filtering now, spam and antivirus filtering. But if I want to compromise the end-user, I am going to send them a URL and redirect them to some website that hosts malware. The stuff that can

become most widespread is SQL injection of some legitimate server, and users will see an iframe that loads a malware onto it.

Experts also commented on the strategic position of the browsers in the fight to protect consumers. First, web browsers can warn users directly and effectively. A recent laboratory study showed that when Firefox 2 presented phishing warnings, none of the users entered sensitive information into phishing websites [25]. This study also recommended changes to Internet Explorer's phishing warnings, and Microsoft has already acted on some of them to improve IE 8's warning mechanism. Second, the browser market is fairly concentrated, with two browsers (Internet Explorer and Firefox) accounting for 95% of the total market [26]. Solutions implemented by these two browsers would provide the majority of users with a defense against phishing.

Finally, experts pointed out that ISPs are in the best position to clean up compromised machines, as described earlier.

We asked experts to comment on and prioritize a set of recommendations for securing the computing environment. Experts ranked the following as top priorities.

Recommendation (R3): OS vendors should continue to secure operating systems by implementing secure coding practices, investing in secure vulnerability patching, and building anti-malware capability directly into the operating systems to enhance default security.

To secure the operating system, experts suggested Microsoft protect the hosts file in Windows XP and earlier versions, as done by some Antivirus software [27], to prevent pharming attacks.

Another way to secure the operating system is by continually updating the system with the latest patches, as a fully patched computer with firewall enabled provides a strong defense against exploit-based malware. However, one of the problems with patching is that distributing a patch provides information to criminals about the security vulnerability that is being patched. Criminals can disassemble the patch and compared it to the code that it replaces. Once a new exploit is known, a malware exploit can be quickly built. It currently takes less than three days – sometimes only a matter of hours – between the time a patch is released and the time a malicious exploit appears. After this short period of time, most computers are still vulnerable. Research and application development into securely delivering patches to computers, possibly using public-key cryptography, would help alleviate the problem [11].

Finally, some experts suggested building anti-virus and anti-malware capability directly into the OS. Experts pointed out that XP service pack 2 has a security center with firewalls enabled and suffers fewer attacks than service pack 1 [28]. These experts also praised Microsoft's effort to distribute malware removal tools and updated malware signatures monthly, and argued that Microsoft should provide some default protection to computer users who do not buy anti-virus software.

Recommendation (R4): Stakeholders should focus on improving the security of web applications, providing support and incentives for fixing applications.

Currently, over 70% of phishing websites are hosted on hacked websites or free hosting sites. Many vulnerabilities for web applications exist (e.g. SQL injection, cross site scripting, remote code execution), making them a tempting target for criminals. Experts suggested a few ways to improve the security of web applications. One expert felt that technical authorities such as CERT or APWG should produce a list of most frequently hacked websites and notify the website operators of their vulnerability.

However, not all website operators have the technical capability or incentives to fix the problem. A recent paper by Moore and Clayton showed that 25% of the hosts used for phishing end up being compromised again within a couple of months [29]. If the compromise is due to a lack of technical capability, then there needs to be a way to provide tools and educational resources to help them secure their web application. On the other hand, if repeated compromises are due to a lack of incentives to fix, then there needs to be a way of punishing transgressors, with escalating consequences.

Another approach is to involve the hosting provider. For example encourage these providers run intrusion detection on the applications they are hosting, and scanning newly created pages for phishing and malware.

Recommendation (R5): Web browser vendors should continue to improve the performance of integrated browser anti-phishing warning systems, with a goal to catch 90% of phishing URLs within an hour after they go online.

As mentioned previously in this section, web browsers is at a strategic position as they can warn users effectively, and faster than other methods. Currently, browser-integrated phishing warning systems catch only 40-60% of the URLs 3 hours after the attacks are launched [30]. To provide the majority of Internet users with adequate protection, these warning systems should be improved.

To accomplish this, the key is heuristics. Currently major browsers only use human-verified blacklists. To raise detection rates significantly, heuristics need to be used to supplement existings blacklists and block attacks more quickly [30]. Another way to improve the coverage of the blacklists is to gather phishing feeds from multiple sources to maximize their coverage [31]. However, as discussed in the next section, browser vendors are extremely cautious in using heuristics because of false positives, incorrectly labeling a legitimate site as phishing, which could potentially expose them to costly lawsuits. We present recommendations to address this issue in the next section.

Recommendation (R6): Academics and for-profit protectors should develop better techniques to quickly identify botnets and proxies, shut down botnet command and control, and clean compromised machines.

To shut down botnets, experts recommended that we either go after their command and control centers or clean the bot machine themselves.

In November 2008, a hosting company named McColo that hosted a bot command and control center was disconnected by

its upstream providers, causing a nearly 70% drop in spam volume [32]. More efforts to identify and shutdown command and control centers would diminish the usefulness of other bots. However, we have to be mindful that criminals will continue to regroup and attack again. A good illustration is that two months after the McColo case, the spam volume was back to the previous level [33]. Spammers find other bot command and control centers, and they are getting more sophisticated in using P2P tools to control bots instead of traditional IRC commands. Defenders need to learn from successes and failures to ensure faster reaction in the future.

The McColo case offers several lessons. There invariably exists some rogue hosting companies (also known as bullet-proof hosting), so persuading them to clean up their network would be difficult and likely have limited effect. Therefore it is important to involve upstream connectivity providers. However, these providers face some challenges for proactive monitoring. For example, the infrastructure for monitoring is expensive, the legal justification is unclear, and because of contractual agreements, they are likely to be very cautious. So other stakeholders such as public protectors or for-profit companies needs to help provide as much evidence as possible. Second, media can play an important role. In the case of McColo, a Washington Post report played a critical role in persuading the upstream providers. Similarly, the media played an important role in having the Russian authorities shut down the Russian business network, a known hosting provider for Internet miscreants [34]. Finally, the higher the level of coordination between stakeholders, the better they are at identifying and shutting down these rogue providers.

Another approach focuses on cleaning up individual machines. This is a much more challenging task as there are millions of compromised machines to fix. ISPs need to be involved. Recognizing the disincentives mentioned in section VI, one expert suggested a notice and take down approach: certain third parties can notify an ISP that a certain computer on its network is in a botnet or doing something malicious. Once the ISP receives the notification, it becomes obligated to clean up the machine. Some experts also pointed out that the U.S government can apply some of the technologies used in national defense to detect malicious activity and clean compromised machine to fix ISP networks.

However, efforts are needed to automate the clean up process. Experts suggested that we won't see much of an impact on crime rate until we clean up a large fraction of compromised machines. Hence, better automatic solutions are needed to complement the notice and take-down approach.

Although no actions have been taken so far, the ISPs we interviewed acknowledged that compromised machines are a big problem. During the interviews, they asked about academic research on automated tools to quarantine these compromised machines. We suggest conducting more research and development focusing on automated mitigation of malware-infected computers.

B. Organizations are conservative about filtering and warning about phish because they are worried about false positives. However, this often leads to repeated efforts and slow reaction.

The issue of false positives came up frequently during our interviews. Generally speaking, phishing detection falls into two categories: blacklist-based methods in which humans verify suspicious phishing URLs, and heuristic approaches that utilize HTML or content signatures to identify phish automatically. In our interviews, we found that providers favor blacklists over heuristics, and even those who do use heuristics are using them conservatively. For example, an expert at an ISP told us that they had a system that warns users if a certain email appears to be phish (based on blacklists and heuristics), but they did not delete these emails because they consider their false-positive rate to be too high.

Browser vendors are also extremely concerned about false positives. The expert from a major browser vendor said that they take false positives very seriously and manually verify each URL on their blacklist to avoid false positives. All of the major browsers appear to favor human-verified blacklists with extremely low false positives over heuristics that may potentially have higher false positives.

Registries consider false positives as their biggest concern in implementing anti-abuse policies. One registry told us that they do not take act on phishing URLs submitted by third parties (such as takedown vendors) until the URLs have undergone a review process to determine if they are really phishing URLs. In other words, a phishing site is verified multiple times by different parties before action is taken, wasting precious time.

Infrastructure providers are concerned about potential liability from mislabeling or taking down legitimate websites. There have been cases where companies have attempted to hold service providers responsible for false positives, but as of yet no company has been held responsible. For example, in a 2005 court case, Associated Bank-Corp sued Earthlink after the Earthlink anti-phishing software ScamBlocker blocked the bank's legitimate page [35]. Earthlink was able to fend off the suit on the basis that it was using a blacklist of phish provided by a third party, thus, under a provision in the Communication Decency Act (CDA), it could not be held liable as a publisher when that information is erroneous. Although the bank apparently did not sue the provider of the blacklist, the court opened the door for them to do that.

False positives based on heuristics have more subtle concerns. If heuristic-based software blocks a phish that turns out to be a false positive, the vendor may be regarded as a publisher under the CDA, and thus not immunized. Because of these fears, heuristics are not favored in integrated browser phishing protection.

It is unclear, however, how future cases, if any, will be handled. One legal expert thought there was no case to be made. He said:

I think everything will depend on what statements are made about the blocked site by the anti-phishing software. For example, when it says, 'we think this site might be a phishing site,' unless they were grossly

negligent (in which case the thinking would not be reasonable), there would probably be no liability. If it said 'This site is absolutely a phishing site' it would be a whole different story.

It is worth noting that vendors have developed blacklist processes and heuristics with extremely low false positive rates. One software vendor told us at their current false positive rate is so low that a user would encounter a false positive only once in a few years. Another takedown provider told us that they only had one or two false positives in the past four or five years, and even those false positives were arguably true positives. Recent academic work has shown that heuristics seem to detect websites with near zero false positives ([36], [30]). It is therefore, unclear why vendors remain so reluctant to use heuristics more aggressively.

To address this issue, we introduce three recommendations based on our findings.

Recommendation (R7): Clarify the legal issues surrounding false positives of blacklists and heuristics.

Companies are adopting conservative strategies to avoid false positives for fear of liability, even when false positives occur rarely. This is hurting phishing protection, especially when heuristics offer real-time protection against phishing and have considerable benefits over blacklists. We encourage more discussion on liability surrounding the use of phishing blacklists and heuristics. So far, there has been no test case on this matter. The question at hand is at what level of accuracy heuristics can be applied to block phish and not be held liable? Some experts argued that zero false positive is the only acceptable level, but most of the experts interviewed feel that it would be reasonable to block with less-than perfect accuracy if a procedure were in place to correct errors. Safe harbor legislation, which immunizes providers from liability if they meet certain standards, may be necessary to make companies comfortable that they will not be held liable.

Clarifying liability is important because lack of clarity on these matters could further reduce vendors' incentives to use heuristics to detect phishing and get protections in place rapidly. Major browser vendors and ISPs potentially take on liability for false positives, but do not lose money directly from phishing. Therefore, an uncertain legal situation may reduce their willingness to be proactive.

Recommendation (R8): Create a central clearinghouse to quickly verify phishing reports coming into APWG and on vendor blacklists.

Currently there is a great deal of duplicated effort as phishing reports end up getting verified by multiple sources. For example, many vendors and service providers will not trust phishing reports until they have verified them themselves. A verification organization could serve as a clearinghouse for phishing reports and allow these reports to be verified rapidly using a standard process in which the evidence supporting each report is fully documented. In addition, it is important to report whether each phishing site is a domain setup for phishing or a legitimate domain that has been hacked. This distinction is important for registrars and

registries, as these cases require different actions to be taken.

Recommendation (R9): Researchers should focus on heuristics that minimize false positives.

A sampling of published research has found that current anti-phishing heuristics have a false positive rate of 0.43% - 12% [30]. However, to make sure these heuristics are used, the false positive rate needs to be extremely low. Since billions of websites are visited each day, even if a heuristic has a 1% false positive rate, it means millions of webpages are falsely labeled. For heuristics to be used widely, the false positive of heuristics needs to be at near zero levels. Recent efforts such as [39] and [40] is a good start.

C. Registrars and registries can play an important role in fighting phishing.

As mentioned earlier, registrars and registries have been generally regarded as lagging in terms of phishing countermeasures, but many experts interviewed agreed that they could play a more active role. For example in the case of fast flux attacks, registrars need to be prepared to suspend phishing domains. The Anti-Phishing Working Group produced a set of recommendations for registrars and registries [15].

One key player is the Internet Corporation of Assigned Names and Numbers (ICANN). It is responsible for managing the root zone DNS, setting and negotiating contractual standards for registrars and registries. ICANN is not a regulatory body like the Federal Communication Commission (FCC) and it has limited capabilities to regulate. Going forward, many experts think that ICANN can and should play a more active role in combating phishing and other crimes. Experts suggested that ICANN establish a minimum set of standards for registrars and registries, coupled with self-regulation and better enforcement. However, experts acknowledged that ICANN needs to play a delicate role and achieve consensus with the parties involved to avoid backlash.

We asked experts to comment on and prioritize a set of recommendations for registrars and registries. Experts ranked the following recommendations as top priorities.

Recommendation (R10): ICANN should improve enforcement of domain abuse.

Experts agree that one thing ICANN can do better is to enforce compliance. One expert familiar with ICANN said:

Some registrars ... are very good at enforcing compliance. Other registrars are very good at looking as if they can't do it. KnujOn lists top 10 registrars with domain abuses. Most of my anecdotal research, we see those same names that come up again and again. But they are just confident enough to keep their accreditation.

ICANN has been improving their efforts. In October 2008, they de-accredited one of the ill-behaving registrars. Experts think more of these efforts would be good, because de-accreditation produces a credible penalty for non-compliance, as it essentially terminated the registrar's business.

Recommendation (R11): ICANN should encourage registries to adopt anti-abuse policies.

Several registries have implemented anti-abuse policies, and anecdotal evidence [37] suggests that registries who have implemented anti-abuse policies have much less fraud than those who have not. An expert who works for a registry that recently adopted anti-abuse policies told us his company adopted these policies after they observed how similar policies helped other registries.

However, some registries may not have enough incentives to adopt anti-abuse policies because adding policies creates overhead. ICANN can provide some incentives. One way to encourage adoption is for registries who have adopted anti-abuse policies to share their stories and explain how they led to cost savings and how they handle the issue of false positives. To some extent this is already being done, but ICANN can encourage this further. Another inducement to adopt anti-abuse policies is for ICANN or APWG to publish phishing takedowns, and to share this information regularly with registrars and registries. Finally, as a stronger incentive, ICANN could use anti-abuse metrics as part of their evaluation criteria for future registry applications, for example approving new gTLDs.

VIII. LAW ENFORCEMENT AND EDUCATION

A. Experts agreed that law enforcement should be emphasized, but law enforcement lacks the necessary tools, personnel, and resources to catch phishers.

Experts agreed that law enforcement is essential to deter phishers, and the top priority for law enforcement anti-phishing efforts is to catch organized phishing operations such as rock phish, which are responsible for more than 50% of the phishing attacks. One expert commented:

If we can take out the major hubs, it is not going to solve the problem, but it can show that law enforcement can catch them ... On top of that, these criminals have complex network, and it is not easy to set up. If we can get these gangs, then we may still have the coding kiddies, but those are a lot easier to catch.

However, experts acknowledged that law enforcement face significant challenges:

International nature of the problem. Experts acknowledged that the underground economy is very specialized. One gang is using compromised web servers in many countries that launch attacks with victims in multiple countries. Currently the Mutual Legal Assistance Treaty (MLAT) forms the basis for cooperation between different nations. However, the law enforcement experts that we interviewed complained that this process is very slow.

Proxies. Phishers use proxies so that it is difficult to catch them when they check balances on compromised accounts. This problem is hard to overcome, as there are estimated to be over 10,000 active proxies and it is necessary for law enforcement agents to perform network monitoring of the

proxy machine to catch phishers. However, a warrant is required for law enforcement to legally monitor proxy machines, and by the time a warrant has been issued, the phisher has moved on to a different proxy.

Lack of accuracy in Whois data: Phishes are aware that law enforcement uses Whois data to trace illegal activity, so phishes fabricate contact information when they register domain names using stolen credit cards.

Lack of analytical capabilities: Law enforcement often lacks the necessary analytical capabilities in determining investigative priorities (including qualified investigators and software tools). One law enforcement officer that we interviewed said:

It takes a lot to identify a criminal. There is a lot of data submitted to us from members of APWG or DPN (Digital PhishNet). We don't have time to look at it all. We have to pick out a few variables we know historically told us that is a good target. But the question is that what are we missing? Is there something on that phishing kit are we missing?

Lack of case development tools to process the subpoena request: Multiple law enforcement agents commented on the lack of case development tools. One local law enforcement agent commented:

When we issue subpoenas, some will give searchable PDFs, others give us Microsoft Access database, and some even give us paper. We need tools to conform to the same form of dataset. This is usually done case by case. If law enforcement has a centralized place to do that so that agents all over the country can use it.

We asked experts to comment on and prioritize a set of recommendations for more effective law enforcement. Experts ranked the following recommendations as top priorities.

Recommendation (R12): Improve and invest more into law enforcement, specifically for international cooperation.

Experts commented that it is currently fairly difficult to cooperate with different law enforcement in different jurisdictions because there is often not a lot of money set aside for cooperation. At this time, the cooperation is through the MLAT process, which is very slow. One way to improve on this is to have a joint-task force between two police jurisdictions.

Recommendation (R13): The US Government should invest in technologies to provide law enforcement with better analytical capabilities to prioritize and manage cases, and expand scholarship program to recruit qualified analysts.

There are over 40,000 classic phishing attempts every month, and prioritizing which cases to pursue is critical. One expert said:

Just speaking on [our organization's] behalf, we get a lot of information in, but we are overloaded. People can share data now, that's occurring, but what's not

happening is the analysis piece. We have limited resources ... We do it manually. We need resources, software and hardware to enable that, also more bodies looking at it. There is no magic about the data, but the magic is in the analysis... taking institutional knowledge and applying some data mining algorithms.

To attract more talent, the US government could provide incentives to recruit from top computer science programs around the country for digital forensics and analysis, possibly through the expansion of DHS scholarship program.

Recommendation (R14): Get more corporations to aggregate and submit fraud data to law enforcement to identify proxies.

Currently, most phishing attacks are from botnets and proxies and almost all criminal organizations use proxies to check account balances of phished accounts. Aggregating these data from various sources will help law enforcement to determine where to request subpoenas for wire taps. One way to do this is by having corporations work together and give law enforcement fraud data with a single list of IP addresses that have checked balances on compromised accounts. Another way is for Internet service providers who have information to share that with law enforcements.

Recommendation (R15): Continue to strengthen collaboration between public protectors, private protectors, and between law enforcement in different countries.

Collaboration is key to catch phishers due to the international nature of phishing. It is vitally important for law enforcement to develop good relationships with their peers in other countries. One notable effort is the Digital PhishNet conferences that NCFTA and Microsoft organize each year. More efforts like these are needed.

B. Experts agree that shutting down money trails is very important to defeat phishers.

Experts said that shutting down the money trail can make phishing less attractive. For example, phishers often use "money mules," persons recruited to receive stolen funds (or goods bought using stolen funds) and then transfer the money out of the country. Mules are recruited by a variety of methods, including spam emails, advertisement on genuine recruitment web sites and newspapers, approaching people who have their CVs available online, and instant messaging.

To shut down money trails, one expert recommended we find out where the mules typically are and how mules are recruited. Another expert suggested that banks and take-down organizations put more effort into shutting down mule recruitment websites. He mentioned recent research that mule recruitment sites takes much longer to shutdown than normal phishing websites.

Another expert proposed a clearinghouse of accounts where each participating bank submit accounts that have been used as mules. Currently, bank fraud systems can detect some suspicious transactions to mule accounts, but there is no

system in place to share this information with other banks. If this list of suspicious accounts were shared, a lot of money laundering could be stopped.

C. Education and awareness are important factors that are not emphasized enough. However, not all experts agree on the effects of education.

Most experts agreed that anti-phishing education for end users needs to be implemented better. However, some experts strongly endorse it, while others say education should *not* be a focus. Both sides have strong words to say. For example, one expert in favor of more education said:

There needs to be some accountability on Internet users People still click on URLs they shouldn't. So we need to stress user education, and a little bit of common sense. We are a society that becomes desensitized to our responsibility. You really end up paying for this over time. You are going to end up paying high interest rates. So you really do need to pay more attention.

Another expert who has worked on anti-phishing campaigns at a large US institution doubted the efficacy of such efforts:

My experience of education is that it won't make that much difference. You have to do it, because if you don't, consumers will get mad at you. There is trust and there is safety. You have to balance both of them. ... However, education doesn't impact phishing losses, or make it less. It doesn't do any of that, what it does is making people feel safer. If your goal is to improve security, then education shouldn't be of top priority."

Based on these comments, we introduced a set of recommendations.

Recommendations (R16): Academic researchers and industry should continue to make education fun, engaging and up to date.

Current academic research shows that popular online user education materials are effective if users *actually* read them. For example, Kumaraguru et. al asked users to read four popular training materials online and tested their ability to recognize phishing websites. They found that users were able to distinguish phishing websites from legitimate ones much better after reading these training materials [40]. However, the problem is that users normally don't read security training materials [41].

To make education more effective, we recommend developing more innovative ways to make education fun, engaging, and up to date (e.g. [42],[43]).

Recommendation (R17): Launch an education campaign to educate the public about mules, and encourage social networking sites to take the initiative to educate their customers about phishing.

Experts mentioned the need to educate money mules, some of whom unknowingly become accomplices to crimes. To educate mules, experts recommend we find out where the

mules typically are and how mules are recruited. Finding out where they are recruited can help determine whether national campaigns or if targeted campaigns are needed.

Experts also thought social networking sites should take the initiative to educate their customers about phishing, as they are increasingly becoming targets of phishing campaigns.

Recommendation (R18): Complement education with other countermeasures such as filtering and better user interfaces.

Where possible, efforts should focus on automatic filtering that does not require user knowledge, and designing better user interfaces that make it more obvious to users what the right trust decision is.

However, education remains an important part of combating phishing because it is unlikely that any automated system will ever be completely accurate in detecting phishing attacks, especially when detection requires knowledge of contextual information. There will still remain many kinds of trust decisions that users must make on their own, usually with limited or no assistance.

IX. CONCLUSIONS

In this paper, we reported on seven findings (summarized in Table 3) and 18 recommendations (summarized in Appendix A) from 31 qualitative interviews with anti-phishing experts.

Our findings suggest that phishing is evolving into a more organized effort. It is part of a larger crime eco-system, where it is increasingly blended with malware and used as a gateway for other attacks.

Experts identified several places where incentives for fighting phishing may be misaligned, in the sense that the stakeholders who are in a position to have the largest impact do not have much incentive to devote resources to anti-phishing. To resolve this, we recommend further study of these misalignments and development of policy alternatives to correct them.

In terms of countermeasures, experts identified improving law enforcement and shutting down money trails as top priorities. We identified key difficulties that law enforcement organizations face, and recommend investment into specific types of technologies made to equip law enforcement to better prioritize cases. Collaboration is the key in these investigations, so we recommend ways to foster it.

Finally, experts agreed that education is an important factor that is not emphasized enough, however, they did not agree on the effects of education. We recommend developing more innovative ways to make education fun, engaging and up to date and propose content areas that education needs to be focused on.

ACKNOWLEDGMENTS

The authors would like to thank the 31 experts who participated in our interviews. This work was supported in part by National Science Foundation under grant CCF-0524189, and by the Army Research Office grant number DAAD19-02-1-0389. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the National Science Foundation or the U.S. government.

REFERENCES

- [1] Messagelabs, "MessageLabs Intelligence May 2009," 2009. [online]. Available: <http://www.messagelabs.com/intelligence.aspx>. [Accessed: Jan. 10, 2009].
- [2] D. Florencio and C. Herley, "A large-scale study of web password habits," In WWW '07: Proceedings of the 16th international conference on World Wide Web, pages 657–666, New York, NY, USA, 2007. ACM Press.
- [3] T. Moore and R. Clayton, "Examining the Impact of Website Take-down on Phishing," In eCrime '07: Proceedings of the 2007 e-Crime Researchers summit, pages 1–13, New York, NY, USA, 2007. ACM Press.
- [4] Gartner Research, "Number of Phishing E-Mails Sent to U.S. Adults Nearly Doubles in Just Two Years," Press Release, 2006. [online]. Available: <http://www.gartner.com/it/page.jsp?id=498245>. [Accessed: Jan. 10, 2009].
- [5] Google Inc, "Google safe browsing for Firefox," 2007. [online]. Available:<http://www.google.com/tools/firefox/safebrowsing/>. [Accessed: Jan. 10, 2009].
- [6] Microsoft Corporation, "Phishing filter: Help protect yourself from online scams," 2008. [online] Available: <http://www.microsoft.com/protect/products/yourself/phishingfilter.msp>. [Accessed: Jan. 10, 2009].
- [7] Apple Inc. Visited Jan 1, 2009. <http://www.apple.com/safari/features.html#security>.
- [8] J. Franklin, V. Paxson, A. Perrig, and S. Savage, "An inquiry into the nature and causes of the wealth of Internet miscreants," In CCS '07: Proceedings of the 14th ACM conference on Computer and communications security (New York, NY, USA, 2007), ACM, pp. 375–388.
- [9] G. Keizer, "Phishers Beat Bank's Two-factor Authentication," Information Week, 2006. [online] Available: <http://www.informationweek.com/news/showArticle.jhtml?articleID=190400362>. [Accessed: Jan 10, 2009]
- [10] Financial Services Technology Consortium, "Understanding and Countering the Phishing Threat," White Paper, 2005. [online]. Available: http://www.fstc.org/projects/docs/FSTC_Counter_Phishing_Project_Whitepaper.pdf. [Accessed: Jan 10, 2009]
- [11] Identity Theft Technology Council, "Online identity theft: Phishing technology, chokepoints and countermeasures," 2005. [online]. Available: <http://www.antiphishing.org/Phishing-dhs-report.pdf>. [Accessed Jan. 10, 2009]
- [12] National Consumers League, "A Call for Action: Report from National Consumers League Anti-Phishing Retreat," Report, 2006. [online]. Available: <http://www.nclnet.org/news/2006/Final%20NC%20Phishing%20Report.pdf>. [Accessed: Jan. 10, 2009].
- [13] Dept. of Homeland Security and the Anti-phishing Working Group, "The crimeware landscape: malware, phishing, identity theft and beyond," Report, 2006. [online]. Available: http://www.antiphishing.org/reports/APWG_CrimewareReport.pdf. [Accessed: Jun. 10, 2009].
- [14] Anti-Phishing Working Group, "What to Do If Your Website Has Been Hacked by Phishers," Report, 2009. [online]. Available: http://www.apwg.org/reports/APWG_WTD_HackedWebsite.pdf. [Accessed May. 18, 2009].
- [15] Anti-Phishing Working Group, "Anti-Phishing Best Practices Recommendations for Registrars," Report, 2008. [online] Available: http://www.apwg.org/reports/APWG_RegistrarBestPractices.pdf. [Accessed May. 18, 2009].
- [16] Message Anti-Abuse Working Group and Anti-Phishing Working Group, "Anti-Phishing Best Practices for ISPs and Mailbox Providers," Report, 2006. [online]. Available: <http://www.apwg.org/reports/bestpracticesforisps.pdf>. [Accessed May. 18, 2009]
- [17] R. Rosenthal, R.L. Rosnow. Essentials of Behavioral Research, third ed. McGraw Hill, New York, NY, USA, 2008.
- [18] S. Dynes, H. Brechbuhl and M. E. Johnson, "Information Security in the Extended Enterprise: Some Initial Results From a Field Study of an Industrial Firm," Fourth Workshop on the Economics of Information Security, Harvard University. Available:<http://infosecon.net/workshop/pdf/51.pdf>. [Accessed: Jun 10, 2009]
- [19] B.R. Rowe and M. P. Gallaher, "Private Sector Cyber Security Investment: An Empirical Analysis," Fifth Workshop on the Economics of Information Security, 2006, Cambridge, UK. Available: <http://weis2006.econinfocsec.org/docs/18.pdf>. [Accessed: Jun 10, 2009]
- [20] T. Moore, R. Clayton and H. Stern, "Temporal Correlations between Spam and Phishing Websites," 2nd USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET '09), April 21, 2009: Boston, MA.
- [21] T.N. Jagatic, N.A. Johnson, M. Jakobsson, and F. Menczer, "Social phishing," Commun. ACM 50, 10 (Oct. 2007), 94-100. DOI=<http://doi.acm.org/10.1145/1290958.1290968>
- [22] Symantec, "Symantec Global Internet Security Threat Report," volume XIV, 2009. [online]. Available: http://eval.symantec.com/mktginfo/enterprise/white_papers/b-whitepaper_internet_security_threat_report_xiv_04-2009.en-us.pdf. [Accessed: Mar. 19, 2009].
- [23] Merchant Risk Council, "Annual e-Commerce Fraud Survey Results," March 3, 2009. [online]. Available: <https://www.merchantriskcouncil.org/index.cfm?fuseaction=Feature.showFeature&FeatureID=119>. [Accessed: June 10, 2009].
- [24] H. Varian, "Managing Online Security Risks," New York Times; New York, NY.; Jun 1, 2000. [online]. Available: <http://people.ischool.berkeley.edu/~hal/people/hal/NYTimes/2000-06-01.html>. [Accessed: Jun 1, 2009].
- [25] S. Egelman, L. F. Cranor, and J. Hong, "You've been warned: an empirical study of the effectiveness of web browser phishing warnings," In CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems, pages 1065–1074, New York, NY, USA, 2008. ACM.
- [26] Net Applications. Inc, "Browser market share Q4 2008," [online]. Available: <http://marketshare.hitslink.com/report.aspx?qprid=0&qpmr=15&qpd=1&qpc=3&qpcal=1&qptimeframe=Q&qpsp=39>. [Accessed: Jan. 10, 2009].
- [27] Zone Alarm, "Smart Defense System", 2004 [online]. Available: http://smartdefense.zonealarm.com/tmpl/AdvisoryArticle?action=detail&aid=20040128sa_000013. [Accessed: Jun. 1 2009].
- [28] Avantgarde, "Time to Live on the Network," Technical Report, 2004. [online]. Available:<http://www.avantgarde.com/xxxxtltn.pdf>. [Accessed: Jun 10, 2009]
- [29] T. Moore and R. Clayton. "Evil Searching: Compromise and Recompromise of Internet Hosts for Phishing," 13th International Conference on Financial Cryptography and Data Security. February 23-26, 2009: Barbados.
- [30] S. Sheng, B. Wardman, G. Warner, L.F. Cranor, J. Hong, and C. Zhang, "An Empirical Analysis of Phishing Blacklists," To appear in 6th Conference in Email and Anti-Spam (Mountain view, CA, July 16 - 17, 2009). CEAS '09.
- [31] T. Moore and R. Clayton, "The Consequence of Non-Cooperation in the Fight Against Phishing," Third APWG eCrime Researchers Summit. October 15-16, 2008: Atlanta, GA.
- [32] B. Krebs, "Host of Internet Spam Groups Is Cut Off," Washington Post, Nov. 12, 2008.
- [33] D. Danchev, "Google: Spam volume for Q1 back to pre-McColo levels," CBS Interactive, April 2, 2009. [online]. Available: <http://blogs.zdnet.com/security/?p=3073&tag=rbxcnbd1>. [Accessed: Jun 1, 2009].
- [34] R.E. Church, "McColo and the Difficulty of Fighting Spam", Internet.com, November 20, 2008. [online]. Available: <http://itmanagement.earthweb.com/features/print.php/3786296>. [Accessed: Jun. 1 2009].
- [35] ASSOCIATEDBANK-CORP v. EARTHLINK, INC. Memorandum and order, 05-c-0233-s 2005. Available:http://www.iplawobserver.com/cases/2005-09-14_Associated_Banc_Corp_CDA_Section_230.pdf. [Accessed: Mar. 18, 2009].
- [36] C. Ludl, S. Mcallister, E. Kirda, and C. Kruegel, "On the effectiveness of techniques to detect phishing sites," In DIMVA '07: Proceedings of the 4th international conference on Detection of Intrusions and Malware, and Vulnerability Assessment, pages 20–39, Berlin, Heidelberg,2007. Springer-Verlag.
- [37] The Anti-phishing Working Group, "Global Phishing Survey: Trends and Domain name use in 2H 2008", May 2009. [online]. Available: http://www.antiphishing.org/reports/APWG_GlobalPhishingSurvey2H2008.pdf. [Accessed: Jun. 1, 2009].

- [38] G. Xiang and J. Hong, "An Adaptive Shingling-based Approach using Search Engines for Zero False Positive Phish Detection," 2009. *Under Submission*.
- [39] B. Pendleton, G. Xiang, and J. Hong, "Augmenting the Crowds: Fighting Phishing on a Budget," 2009. *Under Submission*.
- [40] P. Kumaraguru, S. Sheng., A. Acquisti, L.F.Cranor, and J. Hong, "Teaching Johnny not to fall for phish," Tech. rep. 2007, Carnegie Mellon University. Available:<http://www.cylab.cmu.edu/files/cmucylab07003.pdf>.
- [41] P. Kumaraguru, Y. Rhee, S. Sheng, S. Hasan, A. Acquisti, L.F. Cranor, and J. Hong, "Getting users to pay attention to anti-phishing education: Evaluation of retention and transfer." e-Crime Researchers Summit, Anti-Phishing Working Group, October 4 - 5, 2007, Pittsburgh, USA.
- [42] P. Kumaraguru, J. Cranshaw, A. Acquisti, L. Cranor, J. Hong, M. A. Blair, and T. Pham, "School of phish: A real-word evaluation of anti-phishing training," to appear in Symposium for Usable Privacy and Security (SOUPS), Mountain view, 2009.
- [43] S. Sheng, B. Magnien, P. Kumaraguru, A. Acquisti, L. F. Cranor, J. Hong, and E. Nunge, "Anti-phishing phil: the design and evaluation of a game that teaches people not to fall for phish." In SOUPS '07: Proceedings of the 3rd symposium on Usable privacy and security, pages 88–99, New York, NY, USA, 2007. ACM.

Recommendations	Stakeholders To take action
R1: Financial institutions should produce more accurate estimates of phishing losses and report these statistics.	Financial Institutions
R2: Regulators and academic researchers need to investigate the issue of incentives further.	Regulators and academicians
R3: OS vendors should continue to secure operating systems by implementing secure coding practices, investing in secure vulnerability patching, and building anti-malware capability directly into the operating systems to enhance default security.	Operating System vendors
R4: Stakeholder should focus on improving on the security of web applications.	Various stakeholders
R5: Web browser vendors should continue to improve the performance of integrated browser anti-phishing warning systems, with a goal to catch 85-95% of phishing URLs within an hour after they go online.	Web browser vendors
R6: Academics and for-profit protectors should develop better techniques to quickly identify botnets and proxies, shut down botnet command and control, and clean compromised machines.	Academics, for-profit protectors
R7: Clarify the legal issues of the false positives of blacklists and heuristics.	Legal community
R8: Create a central clearinghouse to quickly verify phishing reports coming into APWG and on vendor blacklists.	APWG
R9: Academics should focus heuristic research on reducing false positives.	Academia
R10: ICANN should improve enforcement on domain abuse.	ICANN
R 11: ICANN should encourage registries to adopt anti-abuse policies.	ICANN
R12: Improve and invest more into law enforcement, specifically for international cooperation.	US Government
R13: US Government should invest in technologies to provide law enforcement with better analytical capability to prioritize and manage cases.	US Government
R14: More corporations aggregating fraud data and submit to law enforcement to identify proxies.	Law enforcement, ISP, academia
R15: Continue to strengthen collaboration between law enforcement in different countries, public and private protectors.	Law enforcement
R16: Academic researchers and industry continue to make education fun and engaging and up to date.	Academician, industry
R17: Launch education campaign to educate the public about mules, and encourage social networking sites to take initiative to educate their customers.	US Government
R18: Complement education with other countermeasures such as filtering and better user interface design.	Academician