Chapter 2

Background

In this Chapter, I will discuss the literature on phishing, relevant countermeasures, and the literature on economics of information security.

2.1 Anatomy of Phishing

Phishing attacks usually take the following steps: planning, setup, attack, collection, fraud and abuse, and post attack [35]. To fight phishing better, we need to understand the nuts and bolts of the attack better. In this section, I described in detail how phishing attacks work. I model in detail both the attackers and the defenders in each of the phishing steps.

Previous works have also modeled phishing attacks. The DHS report on phishing [105] separates phishing into seven steps and discusses countermeasures based on the model. The model does not consider IM phishing and voice over IP phishing, nor does it include stakeholders other than banks. My model includes different stakeholders and provides specific recommendations to counter those two attacks. The FSTC report [35] discusses the phishing life cycle, but does not provide details of attacks, or stakeholders. My analysis addresses these problems.

2.1.1 Planning

Phishers first need to decide whom to target, and what information to steal. A recent study [38] shows that a large and thriving underground ecrime economy of highly specialized criminals exists. Based on this insight, I model phishers as rational agents. They pick and choose their targets to

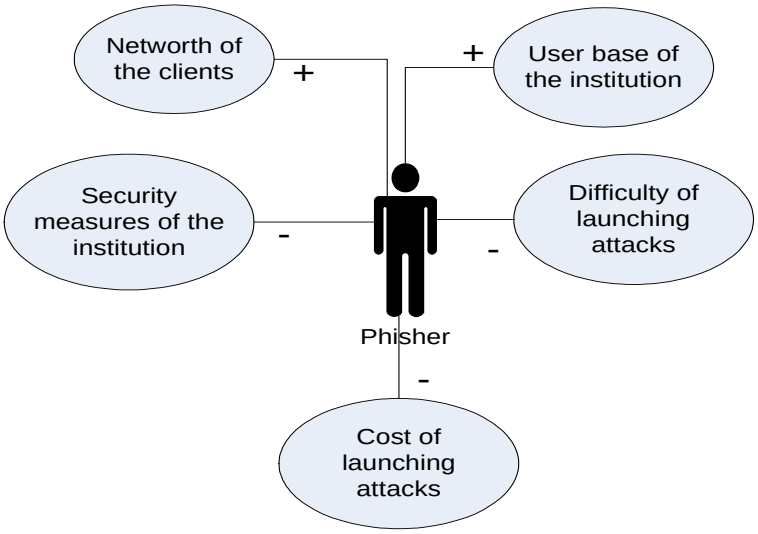

Figure 2.1 A model of a phisher's decision process

maximize their gain and minimize their cost and risks. Figure 2.1 lists the likely factors they consider when they plan the attack.

- **User base of the institution.** The larger the institution's user base, the higher the percentage that targets receiving phishing emails will actually have an account relationship with the institution. Today, each US household carries an average of 6.3 credit cards, and the issuance of these credit cards is concentrated among five banks.[1] Based on my estimate, there is a 15%-45% chance that a household receiving a phishing email from one of these top banks is an actual customer of the bank[2]. This rule does not apply to targeted attacks in which phishers have specific information about customers' account relationships.

- **Net worth of clients.** The higher the worth of the client, the higher the returns of an attack. As countermeasures become more widely available, phishers will launch more attacks against and put high net worth clients such as executives and small business merchant accounts at special risk.

[1] For more information about credit card usage, see FDIC's review at `http://www.fdic.gov/bank/analytical/banking/2005nov/article2.html`

[2] The number of households in US are 110 million in 2010, the top creditcard issue bank (Citibank) have 48 million active cards, the lowest of five (Bank of America) have 20 million accounts.

- **Security measures or processes implemented by the institution.** The stronger an institution's security measures and processes are, the harder it is for phishers to penetrate, commit fraud, and launder money. Phishers can learn about an institution's preparedness through previous attack experience and publicly available information (news and press releases).

- **The value of credentials.** The more valuable the credentials are, the more frequent the attacks will be. Currently small business accounts and checking and money market accounts are for sale at a high price in the Internet black market, whereas credit card numbers are cheap. This means that attacks to steal these accounts will continue to grow. As time goes by, credentials such as social networking data may become more valuable as phishers team up with malware writers to deliver malware and launch phishing.

This analysis of phishers motives and their operating environments yields a few important insights:

- As the underground black market continues to develop, phishing will become an operation that involves multiple parties with different specialties. We can further predict that the economics of scope and scale will be in effect, where phishing operations consolidate into a few phishing gangs to increase profit and reduce cost. This is both good and bad news for law enforcement. The good news is that there are fewer phishers to catch. The bad news is that these gangsters will be more technically capable and advanced enough to hide their trails.

- If a bank is robbed once, we blame the robbers. However, if the same bank is repeatedly robbed, there must be some problem with its security. Simply put, if a particular institution is the repeat target of phishers, its institutional risk control and methods for handling incidents needs to be scrutinized.

- Different categories of targets face different risks and therefore would require different countermeasures.

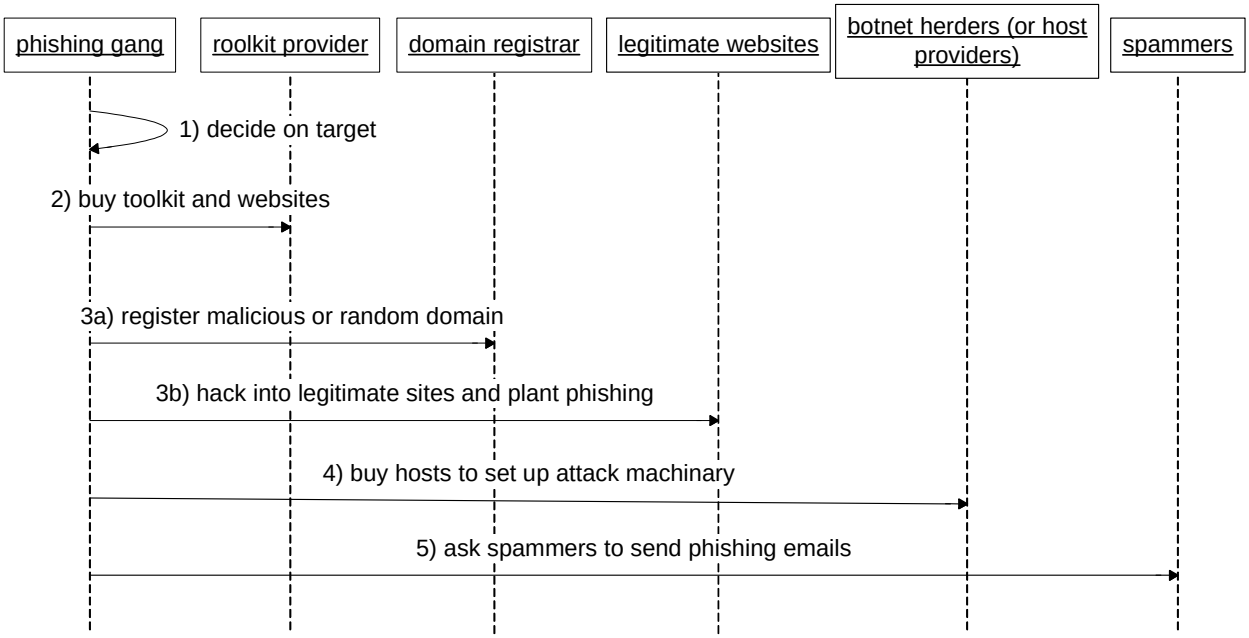

Figure 2.2 Details of Phishing attack planing and setup including stakeholders

2.1.2 Setup

After deciding on the target, phishers will set up the attack infrastructure. Figure 2.2 models how phishers do this and how various stakeholders are involved. The key insight here is that phishers rely on hackers, spammers, botnet herders, and pay loaders to launch a large-scale phishing attack.

A few other insights:

- Phishers will go to great lengths to reduce the probability of being caught. To that end, they will systematically exploit registrars that have weak security or process loopholes (for example, recent rock phish gang exploits of the .hk domain which has a weak verification process), operate from countries that have inadequate law enforcement resources and laws, and deploy proxies to hide their true destinations.

- As shown in Figure 2.2, domain registrars are the first line of contact for phishers. If the registrars improve their security process for registration and fraud detection (for example,

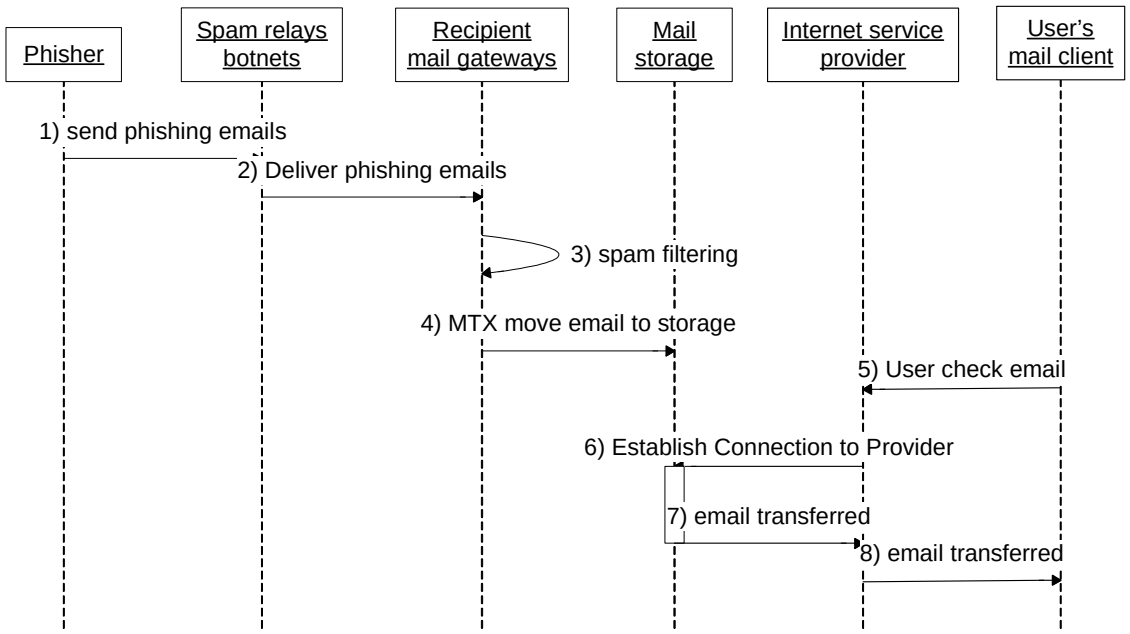

Figure 2.3 Phishing attack via email

simple check against rock phish-type registration), it would sever phishers abilities to get domain names.

- Botnets are the crucial machinery for launching and covering-up phishing attacks. Bots configured as proxies hide the trails of phishing attacks and make it very difficult for law enforcement to investigate.

- All of the criminals meet and trade at the Internet black market. This means that efforts and research to disrupt the Internet black market will not only reduce phishing, but other types of attacks as well.

2.1.3 Attack

Once the machinery is set up, attacks are launched through various vectors. Email and website phishing are the two most common attack vectors. Attacks using instant messenger and voice over IP are also increasing. In this section, I analyze these attacks separately.

Vector: email

Figure 2.3 details the steps of attack through the email and web vectors. First, spammers send phishing emails through spam relays, botnets, anonymous mailers, and other common spam techniques (for a detailed treatment of these techniques, see [142]). The packets arrive at the ISP's mail gateway where emails are put back together. The gateway then performs filtering. After that, mail transfer agents move emails to storage. When users check their email, mail clients on user machines connect to the mail storage through user ISPs, and download emails to their personal machines. If the end user's mail client is a web client, it connects directly to the mail storage through a web interface and retrieves the email.

Figure 2.4 and 2.5 shows an example of phishing email and website from eBay.

Mail providers are in a unique position to combat phishing. They are the first point of contact for phishing emails. Their filtering effort will reduce the magnitude of the problem later on. Their reporting effort will reduce the time for blocking phishing websites. However, we need to consider that mail providers' primary worry is spam. It consumes their bandwidth, takes up server space, and annoys customers. The amount of phishing they receive is only about 1% of the overall spam, so it is likely that they will consider phishing as part of the spam problem. To this end, they would be happy to process phishing if the spam filters also catch them, but they may not be willing to add additional phishing filters that consume their resources. Similarly, they also lack incentive to report phishing emails when it would take manual work to do so.

There are over one thousand commercial free mail providers[3] around the world and many, many more corporate or educational institution email providers. Not all of them are equally resourceful. Major mail providers, such as Yahoo and Microsoft, have resources to invest heavily into anti-phishing and anti-spam technology. Smaller providers with limited IT budgets may have to rely on open source filters such as spam assassin, which is not very effective when it comes to phishing.[4]

[3] Statistics compiled from free email providers' guide [39]

[4] Research has shown that the standard configuration of spam assassin currently only catches about 70% of phishing emails [34].

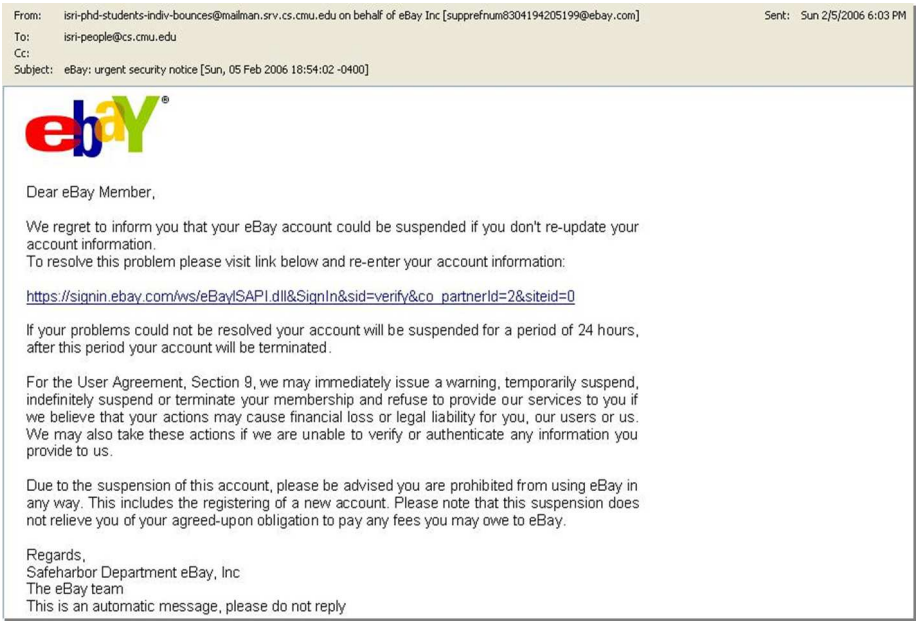

From: isri-phd-students-indiv-bounces@mailman.srv.cs.cmu.edu on behalf of eBay Inc [supprefnum8304194205199@ebay.com] Sent: Sun 2/5/2006 6:03 PM
To: isri-people@cs.cmu.edu
Cc:
Subject: eBay: urgent security notice [Sun, 05 Feb 2006 18:54:02 -0400]

Dear eBay Member,

We regret to inform you that your eBay account could be suspended if you don't re-update your account information.
To resolve this problem please visit link below and re-enter your account information:

https://signin.ebay.com/ws/eBayISAPI.dll&SignIn&sid=verify&co_partnerId=2&siteid=0

If your problems could not be resolved your account will be suspended for a period of 24 hours, after this period your account will be terminated.

For the User Agreement, Section 9, we may immediately issue a warning, temporarily suspend, indefinitely suspend or terminate your membership and refuse to provide our services to you if we believe that your actions may cause financial loss or legal liability for you, our users or us. We may also take these actions if we are unable to verify or authenticate any information you provide to us.

Due to the suspension of this account, please be advised you are prohibited from using eBay in any way. This includes the registering of a new account. Please note that this suspension does not relieve you of your agreed-upon obligation to pay any fees you may owe to eBay.

Regards,
Safeharbor Department eBay, Inc
The eBay team
This is an automatic message, please do not reply

Figure 2.4 An example phishing email from eBay asking users to login to update an account. It warns users that failure to comply will lead to account suspension. The email address is spoofed, and the URL link is spoofed as well.

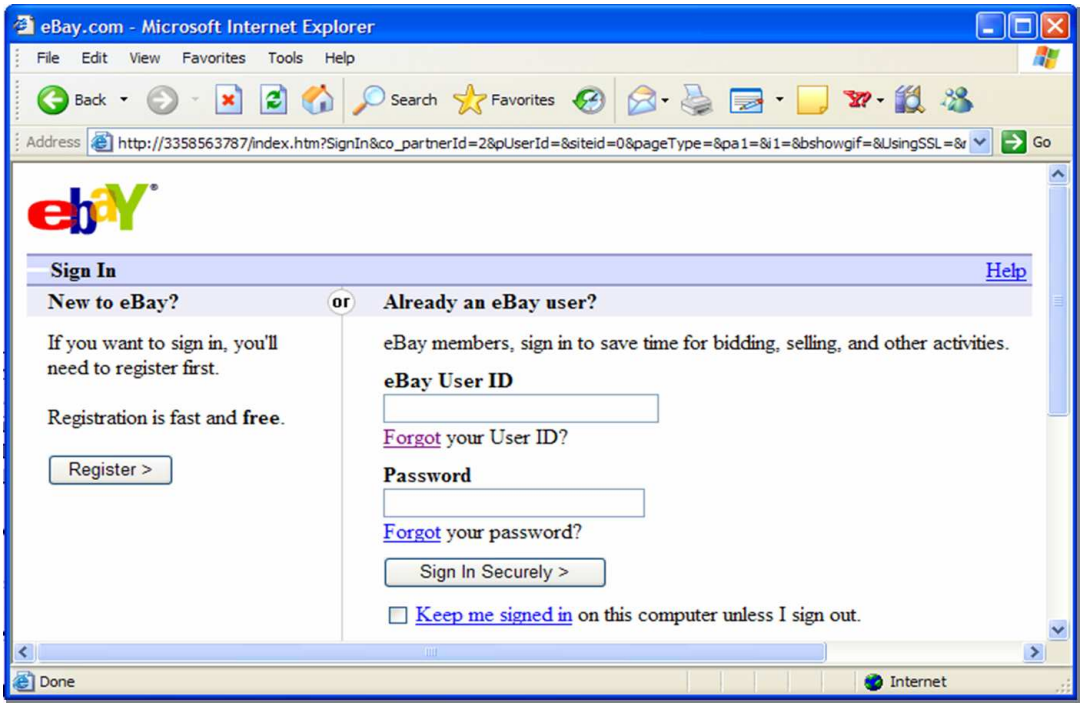

eBay.com - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Back Search Favorites

Address http://3358563787/index.htm?SignIn&co_partnerId=2&pUserId=&siteid=0&pageType=&pa1=&i1=&bshowgif=&UsingSSL=&r Go

Sign In Help

New to eBay? or Already an eBay user?

If you want to sign in, you'll need to register first.

Registration is fast and free.

Register >

eBay members, sign in to save time for bidding, selling, and other activities.

eBay User ID

Forgot your User ID?

Password

Forgot your password?

Sign In Securely >

Keep me signed in on this computer unless I sign out.

Done Internet

Figure 2.5 An example phishing website from eBay that people see once they clicked on the link in the email example above.

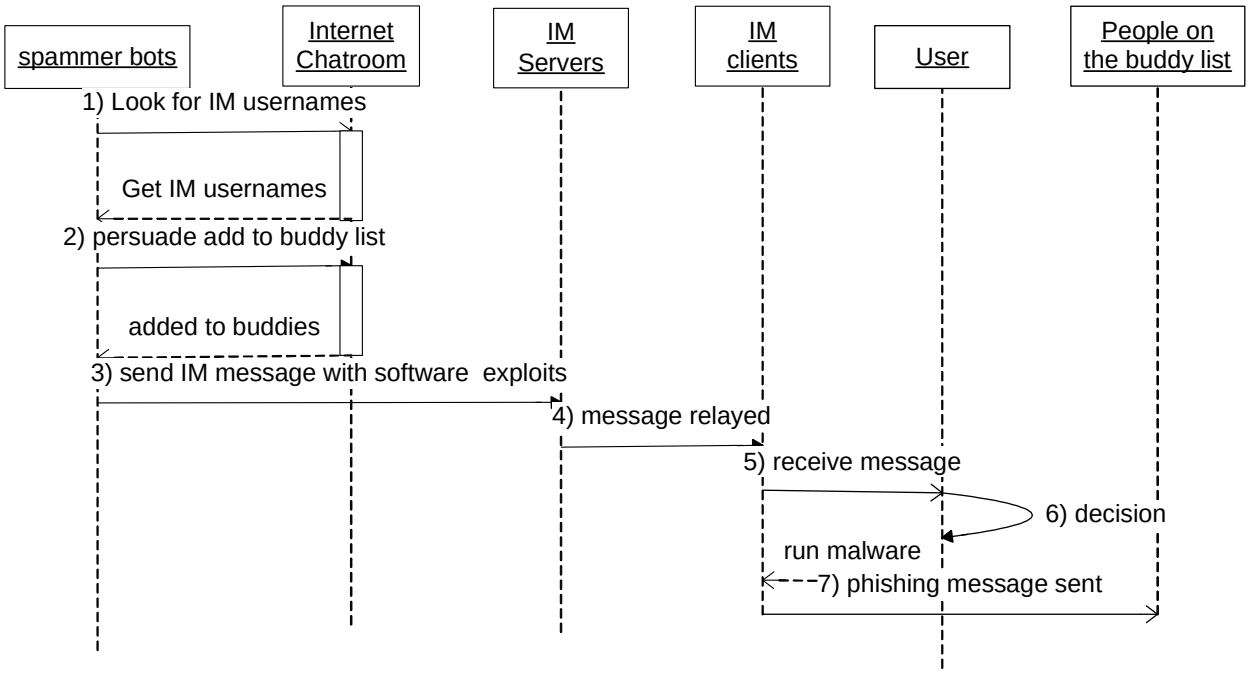

Figure 2.6 Phishing attack via Instant Messenger

Instant messaging phishing

In the instant messaging phishing case, the attack method is very similar to SPIM (spam over IM) and malware. Figure 2.6 documents one particular way to launch the attack. This attack method assumes that not all instant messaging users have configured their clients to receive messages only from their buddy lists. In this example, phishers first get IM names from various chat networks or trick users to add them to their buddy lists. Phishers then send software exploit to hijack account holders' instant messenger lists. If the users run the exploit, phishing messages are sent out as if coming from the users. Alternatively, phishers can also send exploits through email, like a regular malware or virus attack.

Instant messaging phishing and regular email phishing have some similarities and differences:

- In email phishing, users control when they will read messages, whereas in instant messaging, users will generally read message when they are sent. This difference means that IM countermeasures have a shorter window of opportunity to detect and stop messages.

- IM is more contextualized than email. This means that an IM phishing attack targeting a bank will probably not be very effective because users would regard it out of context. However, a message purported from the messaging software vendor (AOL, Yahoo, Microsoft) asking users to update account information would be much more appealing. A message pretending to be sent from a friend to ask to view a picture (while installing a malware that alters the local DNS) could also be highly effective in IM. I expect IM attacks to be more sophisticated and deploy malware-based phishing rather than just regular phishing.

- IM networks are more connected than email networks. If IM networks are attacked, potential infection can have a greater ripple effect than email. Some research has shown that theoretically it would only take 30 seconds to infect 500,000 machines that have IM clients [80].

- IM is easier to control than email. In IM, a few companies own the infrastructure and the delivery channel of the IM network. This means that it is easy to implement control measures at the IM gateway level if messages are routed in peer to server mode [49]. Clients can also connect in peer to peer mode, bypassing the server altogether.

Vector: phishing over VoIP

There are two ways to launch phishing attacks with voice over IP. In the first method, phishers create an exact replica of the target company's voice system, obtain an 800 number from a VoIP service provider, and write and send an email messages that instructs recipients to call the 800 number and verify their credentials. Figure 2.7 illustrates this attack.

In the second method, phishers directly dial victims' phone numbers. Phishers use a random digit dialer or other means to acquire lists of phone numbers to dial (similar to spam over VOIP). Once they have a list, they dial numbers using prerecorded messages, usually alerting consumers of fraud and asking them to verify personal information such as social security numbers and bank pin numbers. Consumers may be tricked into giving credentials on the spot, or by calling back.

In my expert interview, I will ask experts about what can be done regarding VoIP phishing.

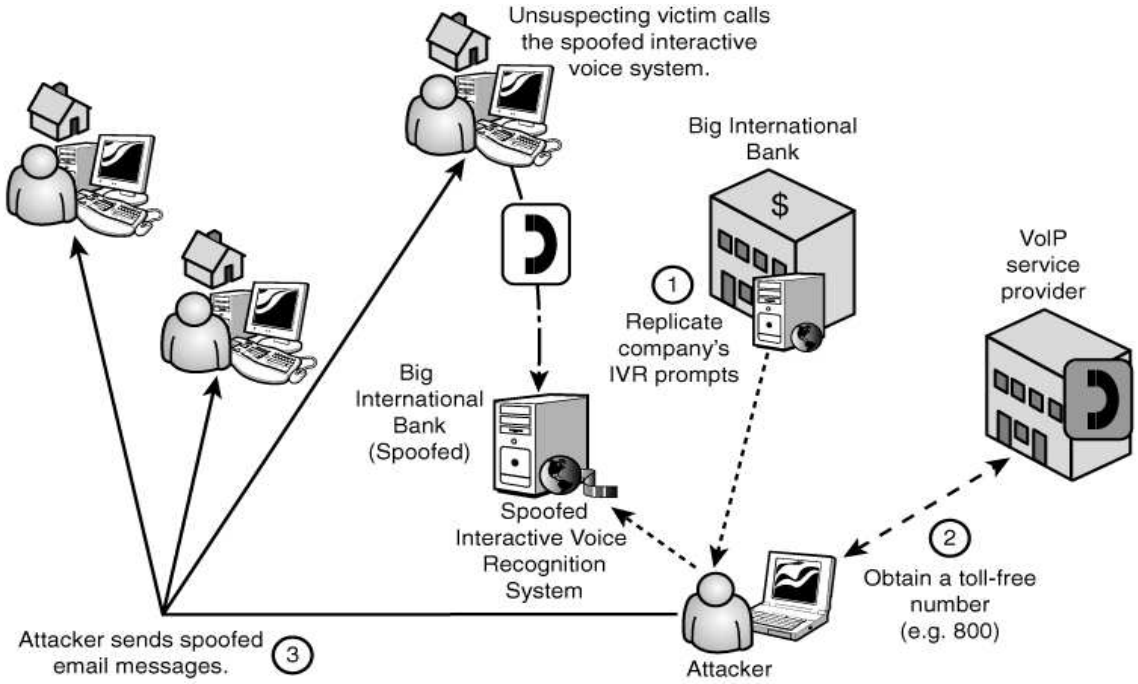

Figure 2.7 Phishing attack via VoIP combined with traditional email scam methods. [132]

2.1.4 Collection

The next step is the collection process. Figure 2.8 shows the process of a user falling for an email-based phishing attack. Other attacks that do not involve the email vector but have the web vector, can also be applied.

Users, email client vendors, browser vendors, and ISPs all have a stake in developing and deploying countermeasures. Below is a list of possible measures.

- Email clients are the first intervention points that can present visible warnings to users. Effective warnings presented here will alert users who were previously unaware of the issue. So far little research exists on the effectiveness of these warnings.

- When users visit phishing websites, most of them have already committed time to read the emails, and it is likely that most will follow the steps of the website and submit information if no warnings were given to them. Studies have suggested that this is indeed the case [68].

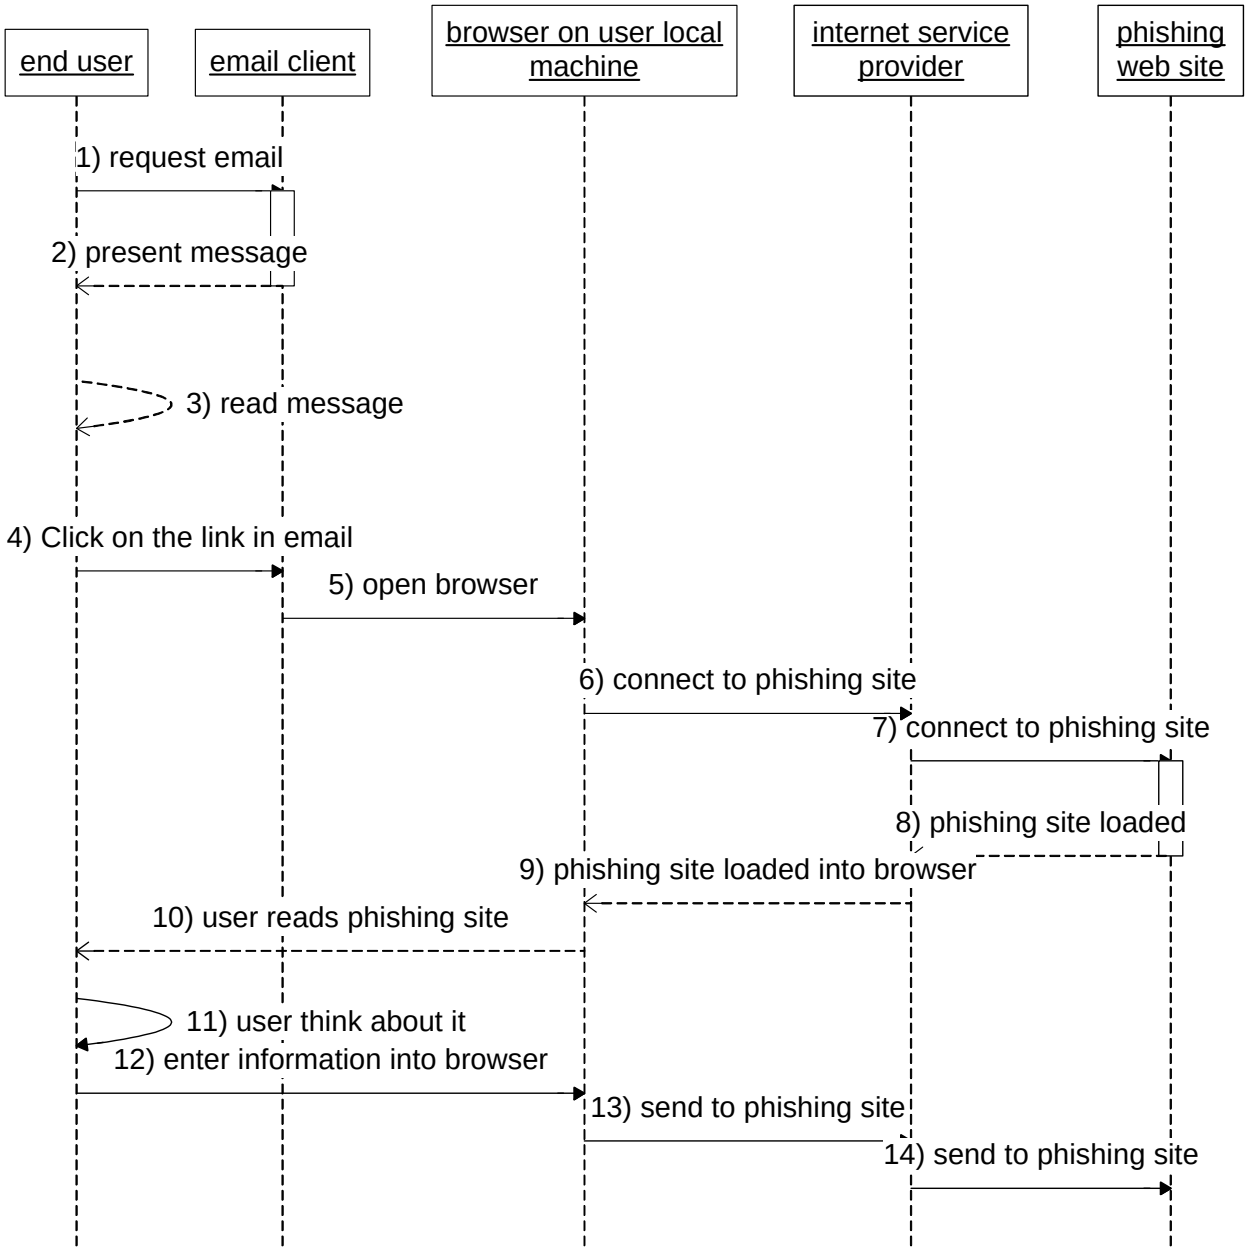

Figure 2.8 Attack and Data collection process

This means that warnings should be focused on earlier in the process, at the email client level. Browser warnings must make it very difficult by design for users to bypass the warnings.

- Browsers are suited to implement solutions, as there are only a few browser products available, compared with tens of thousands of registrars, and ISPs. Currently, major browsers like Internet Explorer and Firefox have phishing filters built into their systems. However, these filters are not so good. For example, in an recent study, Internet Explorer only catches about 50 80% of the phishing urls during the first 6 hours (see figure **??**). Research should be conducted to benchmark and improve its performance.

2.1.5 Fraud

The final process is fraud. Once the information is stolen, it is usually sold on the Internet black market. It is also possible that phishers themselves use the information to defraud customers. The fraud takes place in three steps:

- Phishers use credentials to open new credit cards, or penetrate into user accounts. They would either do this themselves, or employ a cashier.

- Phishers recruit money mules with accounts in the same institution and transfer the money to the mule's account. If they are able to obtain credit cards, they will buy expensive items and have them shipped to the mule.

- Mules take out the money and transfer to the phishers in other countries using Western Union and Money Gram. Phishers may also repackage the expensive goods and ship to a third location.

2.2 Why people fall for phishing

Phishing is part of a larger class of attacks known as semantic attacks. Rather than taking advantage of system vulnerabilities, semantic attacks take advantage of the way humans interact with computers or interpret messages [123]. In the phishing case, attacks exploit the fact that users

tend to trust email messages and web sites based on superficial cues that actually provide little or no meaningful trust information [23, 26].

Research on phishing have shown that people are vulnerable for several reasons. First, people tend to judge a website's legitimacy by its "look and feel", which attackers can easily replicate [23]. Second, many users do not understand or trust the security indicators in web browsers [140]. Third, although some consumers are aware of phishing, they do not link that awareness to their own vulnerability or to strategies for identifying phishing attacks [26]. Fourth, the perceived severity of the consequences of phishing does not predict their behaviors [27]. Below I summarize some of the seminal research in understanding why people fall for phishing.

Dhamija et al showed twenty-two participants twenty web sites and asked them to determine which were fraudulent. Participants made mistakes on the test set 40% of the time. Many of the participants rely on the content of the webpage (logos, layout, graphic design) to determine its legitimacy. The authors noted that 23% of their participants ignored all cues in the browser address bar and status bar as well as all security indicators [23]. Two related study done by Wu et al and Jakobsson showed similar results [54, 140] .

Downs et al have described the results of an interview and role-playing study aimed at understanding why people fall for phishing emails and what cues they look for to avoid such attacks. There were two key findings in their work. First, while some people are aware of phishing, they do not link that awareness to their own vulnerability or to strategies for identifying phishing attacks. Second, while people can protect themselves from familiar risks, people tend to have difficulties generalizing what they know to unfamiliar risks [26].

In a follow up study, Downs et al surveyed 232 computer users to reveal predictors of falling for phishing emails, as well as trusting legitimate emails. Their data suggested that deeper understanding of the web environment, such as being able to correctly interpret URLs and understanding what a lock signifies, is associated with less vulnerability to phishing attacks. However, they also found that the perceived severity of the consequences of phishing does not predict behavior, suggesting that educational efforts should aim to increase users' intuitive understanding, rather than merely warning them about risks [27].

2.3 Cost of phishing

Phishing exerts both direct and indirect cost to the society. Examples of direct loss include consumers losing money, and banking fraud, etc. Examples of indirect cost include erosion of consumer trust of the Internet, negative impact to businesses' brand, , an increase in service call center complaints volume etc. Estimating either cost is hard, as there are many stages of the attack and it is difficult to collect good data. Three reports attempted estimating direct costs.

Gartner Research conducted a survey of 5000 Internet users in August 2006 asking whether consumers have received, clicked or given information in phishing emails. Based on this survey, they estimated that 24.4 million Americans have clicked on a phishing e-mail in 2006, while 3.5 million have given sensitive information. They calculated that the economic loss be 2.8 billion dollars in 2006 [42]. A follow up survey in 2007 with similar methodology estimated that 3.2 billion dollars is lost in 2007 [43].

The above studies rely on people's survey responses. Psychology literature has shown that there is often a wide discrepancy between people's stated choices and their actual behavior.

Moore and Clayton empirically studied phishing websites using PhishTank data. They found that a phishing site lives for 61 hours on average. Using the web log data of some of these phishing sites, they estimated that on average 18 users would fall for phishing on the first day when the site was up, and subsequently 8 users per day afterwards. The total cost to consumers per year was estimated around 320 million dollars [92].

Florencio and Herley [36] instrumented Microsoft's anti-phishing toolbar to send notifications back to Microsoft every time a password was re-used on more than one site. They cross-checked sites that reused password against Microsoft's blacklist, and estimated that 0.4% of people are falling for phishing (the data set had 500k people across a few months). Using this data, and an average of $572 per victim, the estimated yearly cost of phishing to consumers is around 350 million dollars.

The above empirical studies use separate methodologies to estimate phishing costs. Their research designs are reasonable, and their estimates agree remarkably well. We can thus treat 350

million as a lower bound on the direct economic loss of phishing to consumers. The actual cost of phishing exceeds this amount as we consider cost to businesses, and other indirect costs.

It is worth noting that there is disagreement of the magnitude of phishing losses. For example, Florencio and Herley recently used stylized economic models to predict that phishing is a classic example of tragedy of the commons, in which there is open access to a resource that has limited ability to regenerate. Since each phisher independently seeks to maximize his return, the resource is over-grazed and yields far less than it is capable of. The situation stabilizes only when the average phisher is making only as much as he gives up in opportunity cost. They estimate the annual loss to phishing around 60 million dollars [50]. Other security experts have criticized both their research assumptions and results [21]. Such disagreement highlights the lack of empirical evidence for phishing losses.

Last but not least, the cost of phishing is disproportionally borne by consumers. The fraud is relatively small in most phishing scams (medium loss per victim around 200 dollars [43]), but the psychological fear, anxiety it causes the victims, and the time it takes to restore identity can be substantial. However many organizations only consider direct monetary loss to them in their calculation.

2.4 Recent developments in phishing

Traditional phishing is delivered through email, where phishers send mass email to consumers asking them to visit a website to enter information. However, recently attacks have become more sophisticated. In this section, I will talk about three recent developments: voice over IP phishing, spear phishing , and some new phishing techniques such as rock phish and fast flux.

VoIP Phishing

In April 2006, phishers started to use Voice over IP to scam consumers (a.k.a vishing). The attacks work as follows. First, phishers set up a voice-mail system using voice over IP and private branch exchange software (such as open-source PBX software Asterisk). They then use an automatic dialer to call a long list of people and play a recorded message, or simply send emails

asking them to call a number to update their account [112]. When consumers respond, they hear an automated message asking them to enter their account information [138]. Using VOIP, phishers can achive economics of scale by dialing through a long list of numbers, it also makes them harder to track down than using regular phones.

Vishing is growing. MessageLabs has observed the increase frequency of such attacks toward the end of 2007 [86]. Vishing is more damaging than other phishing methods because research has shown that customers generally trust the phone channel more than the email channel [54], and they are well accustomed to enter credit card numbers through automated systems. In my proposed work, I will elicit experts' opinion on how to better counter vishing attacks.

Spear Phishing

Spear phishing is a highly targeted phishing scam. Instead of sending mass phishing emails to everyone, phishers exploit the social context to send spoofed emails to consumers that appear to come from someone they know. For example, phishers can send emails impersonating an organization's system administrator asking people to update their passwords, or impersonate one of your friends in social networking sites. Attacks like these are possible because a myriad of information about consumers exists on the Internet and in many cases are readily available through basic mining of the web. For example attackers can obtain information about consumers' bidding history or shopping preference from eBay, what banks they use (discoverable through web browser history [56]) or even social security number [53].

A recent incident of spearphishing targeted wall street executives. Phishers sent emails to middle and upper senior management of some wall street firms. These emails appeared to be complaints from the Better Business Bureau, and it contained an a .doc attachment where a spying trojan was embedded [86], many opened the attachment and later found they became victims of fraud.

Research studies have shown that spear phishing can be highly effective. In a recent study by Jagatic et al at Indiana University, they send email to students impersonating their friends in the social networking site. The email asked them to enter their secure University credentials. They

found that 72% of the time, users would enter their correct credentials, four times more effective over the traditional phishing methods [53]. Studies at West Point Military Academy showed similar results [33].

In my proposed work, I will also elicit experts' opinion on how to counter spearphishing attacks.

Rock phish and fast-flux

Rock phish refers to phishing committed by the rock phish gang. They were referred to as the Rock phish gang because early versions of their attacks contained the word rock. The rock phish gang has employed several techniques that make them more difficult to defeat than other phishers. First they use stolen credentials to register multiple short and random domain names at multiple registrars. They then host their own DNS servers, and provide name-to-IP service for each of the fraudulently registered domain. The name-to-IP matches to a farm of compromised computers, which do not host the phishing site, but merely act as a proxy to a handful of servers that host phishing sites [81].

Techniques like these pose challenges. Random domain registrations make it hard for automatic detection. Layers of redundancy makes it hard to shut them down quickly, especially if it is from different jurisdictions. A recent study by Moore and Clayton showed that rock phish domains last almost three times longer than regular phishing domains [92].

Another technique the rock phish gang use is fast flux. It is a DNS technique to evade blacklists. It works the following way: multiple nodes within a network register and deregister their address to a single DNS record. Some of these registration lasts only a few hours. Techniques like these pose challenges to blacklists as phishers cycle through hundreds of addresses in a day.

In summary, phishers continue to improve themselves using new techniques such as rock phish and fast-flux, targetting specific groups, and using alternative channels to attack. In my proposed work, I will elicit the experts' recommendations and stakeholder's countermeasures addressing specifically these three new threats, and other evolving threats.

2.5 Phishing countermeasures

Phishing countermeasures can be categorized as legal solutions, technology countermeasures, and social responses. In this section, I will survey each of these types of solutions briefly.

2.5.1 Legal solutions

In the wake of increasing publicity about phishing scams, both federal and state legislatures acted. In January 2005, Virgina added phishing to its computer crimes act, categorizing the use of a computer to obtain personal information "through the use of material artifice, trickery or deception" a Class 6 felony punishable by prison sentences of up to five years and fines of up to $2,500 [137]. Similar statues has been enacted in New Mexico [103], and New York [130]. By February 2007, half of the 50 U.S states have enacted laws addressing phishing [98]. At the federal level, anti-phishing laws have been passed in the House of Representatives in 2004 [133], 2005 [134] and 2007 [135], but the Senate failed to act upon it.

Some prosecutions has been made. In 2006, a Florida man has been indicted in Pennsylvania for a phishing scam that mimicked a Hurrican Katrina relief website [73]. In 2004, Zachary Keith Hill plead guilty in a Texas federal court to crimes related to phishing activity and was sentenced to 46 months imprisonment [45]. The U.S. Department of Justice has successfully prosecuted several other defendants in U.S. courts [118].

However, criminal law does a poor job of deterring phishing because phishers are so hard to find [14]. Law enforcement authorities have little time to track down the criminal through the fraudulent site because on average they only live a few days. Once the site is shut down, the e-mail is the only remaining evidence, and phishers often cover their tracks using such tools as anonymous remailers [14].

In light of this, some legal experts argue that Internet service providers be made liable for part of the Internet "insecurity". Their argument is that "Internet service providers control the gateway through which Internet pests enter and reenter the public computer system. They should therefore bear some responsibility for stopping these pests before they spread and for helping to identify

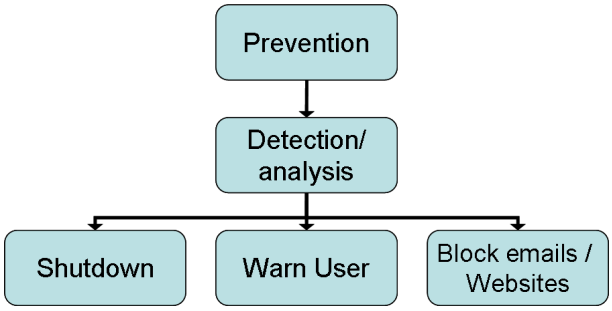

Figure 2.9 Taxonomy of phishing technical countermeasures

individuals who originate malicious code in the first place" [74]. Service providers, however, have largely been immune to such liability. Because they are distributors of the content, not publishers of the content, and as long as they exercise due diligence to remove these materials, they are be not held liable.

How can law enforcements be more effective in the fight against phishing? In my proposed work, I plan to address this issue with further analysis of laws and cases, and by interviewing several law enforcement experts from the Department of Justice and Federal Bureau of Investigation.

2.5.2 Technology countermeasures

Anti-phishing services are now provided by Internet service providers (ISPs), built into mail servers and clients, and available as web browser toolbars. In this section, I will review some commercial offerings as well as academic research. Drawing from the life cycle of phishing attacks, we can categorize countermeasures into the following stages: prevention, detection/analysis, shutdown, block emails/websites and warning users (see Figure 2.9). I discuss each of these stages briefly below.

- **Prevention:** As shown in Figure 2.9, the first step to fight phishing is to prevent attacks before they are materialized. Law enforcement officers can catch and prosecute phishers before they launch the attack. Registrars can monitor domain registrations and analyze suspicious registrations. Mail providers can use email verification solutions such as SPF to drop

Table 2.1 Summary of commercial phishing countermeasures

Stages	**Techniques Used**	**Examples of Companies offering the Service**	**Metrics for Effectiveness**
Prevention	1) monitor domain registrars for suspicious registrations; 2) register domain names defensively to protect a brand; 3) Sender Policy Framework or similar technologies to validate email senders; 4) email encryption using S/Mime or PGP	MarkMonitor, Brandimensions, Cyveillance, InternetIdentity.com, GoDaddy.com, Verisign, TimerWeed Communications, RSA	number of criminals caught and prosecuted, number of phishing attacks stopped
Detection / analysis	1) Set up honeynet or spam traps to collect phishing emails; 2) Scan mail provider's incoming mails; 3) Scan through company weblogs for suspicious activities; 4) User report phishing scams; 5)Scan the web to find malicious websites	MarkMonitor, RSA, Cyveillance	detection time, true positives vs false positives
Block emails / websites	1) blacklist; 2) heuristics	internet service providers	true positives and false positives
Warn User	1) Email client warning; 2) Browser antiphishing toolbars	Microsoft, Google, CloudMark, Earthlink	reach (market share), time to warn, true positives and false positives
Shutdown	Contact ISPs, CERTs or necessary authorities to shutdown malicious website	RSA	time to shutdown, cost of shutdown
Authentication and fraud detection	1) Extended Validation Certificates (EV Certs) 2) Two factor authentication (smart card, tokens) 3) fraud detection system	Verisign, GlobalSign, RSA, Tricipher	To be determined

unverified traffic. The more effective the prevention is, the smaller the phishing problem will be.

- **Detection:** Once phishing attacks are launched, the best defense is to detect and analyze them as early as possible. Internet service providers can add detection systems in their e-mail processing and storage systems to detect suspicious emails. Anti-phishing tool providers can set up spam traps or honeynets to receive early notice of new waves of attacks. Once suspicious emails and websites are identified, analysis will follow, usually combing automatic analysis with human expertise. Table 2.2 lists some of the current state of the art techniques used in automated detection.

- **Shutdown:** Once attacks are verified, service providers can be contacted to shutdown websites.

- **Block emails / websites:** Shutting down websites may take a few days, especially if they are on foreign domains. However, mail providers can delete phishing emails from storage (or move them to a separate spam or phishing folder). Internet service providers can block their customers access to these websites, and replace them with generic education messages.

- **Warning users:** Browsers and, email clients are in a unique position to warn users because their warnings are most visible and direct.

- **Authentication and fraud detection:** This is the last line of defense. Correctly implemented two factor authentication systems can stop phishers from defraud financial institutions, fraud detection system can also discover the scam and stop it.

In summary, the objective is to prevent phishing attacks as much as possible; to detect attacks as early as possible; to shutdown operations as quickly as possible; and to warn users as effectively as possible. Table 2.1 summarizes some of the commercial offerings by the stages, and Table 2.2 summarizes major academic contributions to detection.

In my proposed work, I plan to gain more understanding of the effectiveness of these solutions. Specifically, I will test ten anti-phishing toolbars empirically.

Table 2.2 Meta analysis of proposed detection methods by academics

Authors	Vectors Addressed	Vectors Addressed	Test	Results
Fette et al [34]	Emails	Machine Learning: Random Forests Approach. 6 features	860 phishing emails, and 6950 non-phishing emails	True positive: 96.1%, false positive: 0.1%
Abu-Nimeh et al [1]	Emails	Compared multiple ML techniques, large feature set	1171 phishing emails and 1718 legitimate emails	logistic regression lowest false positive 4.89%, random forest lowest false negative: 11.12%
Yue et al [148]	Websites	Content based approach on TD-IDF with five heuristics	100 phishing and 100 legitimate URLs	True positive: 97%, false positive 10%
Anthony et al [40]	Websites	Use a linear programming model to assess the visual features' similarity	1000 real webpage, 9 phishing page	precision: 99.87%, recall: 88.88%

Figure 2.10 Active warning toolbars in Mozilla Firefox blocking a known phishing site

2.5.3 Social response: awareness and education

Despite claims by some security and usability experts that user education about security does not work [31], there is evidence that well designed user security education can be effective [68].

Web-based training materials, contextual training, and embedded training have all been shown to improve users' ability to avoid phishing attacks.

A number of organizations have developed online training materials to educate users about phishing [28, 32]. In a previous study, we tested the effectiveness of some of these online materials and found that, while these materials could be improved, they are surprisingly effective when users actually read them [70].

Several studies have adopted a contextual training approach in which users are sent simulated phishing emails by the experimenters to test users' vulnerability to phishing attacks. At the end of the study, users are given materials that inform them about phishing attacks. This approach has been used in studies involving Indiana University students [53], West Point cadets [33], and New York State employees [104]. In the New York State study, employees who were sent the simulated phishing emails and follow-up notification were better able to avoid subsequent phishing attacks than those who were given a pamphlet containing information on how to combat phishing.

A related approach, called embedded training, teaches users about phishing during their regular use of email. In a previous laboratory experiment to evaluate our prototype embedded training system, we asked our participants to role play and respond to the messages in an email inbox that included two training emails designed to look like phishing emails. If a participant clicked on a link in a training email, we immediately presented an intervention designed to train them not to fall for phishing attacks. We created several intervention designs based on learning sciences, and found that our interventions were more effective than standard security notices that companies email to their customers [68]. A follow up study showed that people were able to retain what they learned in the training as well [69].

2.6 Economics of Information Security

To fight phishing, companies have to allocate resources. As companies make decisions as how to best allocate resources, it is important to understand incentives and tradeoffs of these decisions. There is a growing field of literature on economics of information security that broadly addresses some of these questions.

2.6.1 Security investment

In general, security investment studies seek to answer the question: what is the optimal amount of investment for information security for a given company? What are the incentives and disincentives that affect the security investment? There are two approaches to answer this question. The first type is through quantitative models, and the second approach is qualitative studies.

Seminal work in quantitative economic modeling includes works by Gordon and Loeb [47], and by Cavusoglu and Raghunathan [15].

For qualitative studies, Row and Gallaher [117] conducted a series of interviews with large organizations in a variety of sectors. Based on the interview, they derived a conceptual approach for security investment in organizations (see Figure 2.11). The paper made three key observations. First, various internal and external incentives (drivers) affects organizations to adopt countermeasures. Second, some organizations tend to adopt more proactive countermeasures while others more reactive. In my proposed work, I extend their model and study the incentives for different stakeholders in phishing countermeasures.

2.6.2 Security as externality

Externalities can be found when we analyze security investment, as protection often depends on the efforts of many principals.

Anderson and Moore used the following analogy to explain this externality.

> Consider a medieval city. If a main threat is a siege, each family is responsible for maintaining and guarding one stretch of the wall, then the city's security will depend on the efforts of the laziest and most cowardly family. If, however, disputes are settled by single combat between champions, then its security depends on the strength and courage of its most valiant knight. But if wars are a matter of attrition, then it is the sum of all the citizens' efforts that matters [4].

Does successfully combat phishing depends on the efforts of the laziest and most cowardly family (weakest link)? or is it depends on the most valiant knight? or sum of efforts?

Kunreuther and Heal notes that security investments can be strategic complements: An individual taking proactive measures creates positive externalities for others that in turn may discourage their own investment [72]. An example is airline security where airlines may decide not to screen luggage transferred from other carriers that are believed to be careful with security.

For phishing countermeasures, externality is clearly an issue. Banks and online merchants suffer loss because the Internet is not secure. On the other hand, the efforts of Internet service and mail providers may reduce the incentive for banks to invest more in phishing, in other words, make them free ride.

2.6.3 Misaligned incentives

Anderson and Moore indicates that incentive misalignment significantly undermines information security [4]. For example, in the United Kingdom, banks are liable for financial fraud only when it is proven that they are at fault. The the burden of proof is with consumers. Therefore customers complaints are not taken seriously, and this leads to lots of fraud [4].

Would there be misaligned incentives among stakeholders of phishing countermeasures? Would those in the best position to fight the phishing lack incentives to do so?

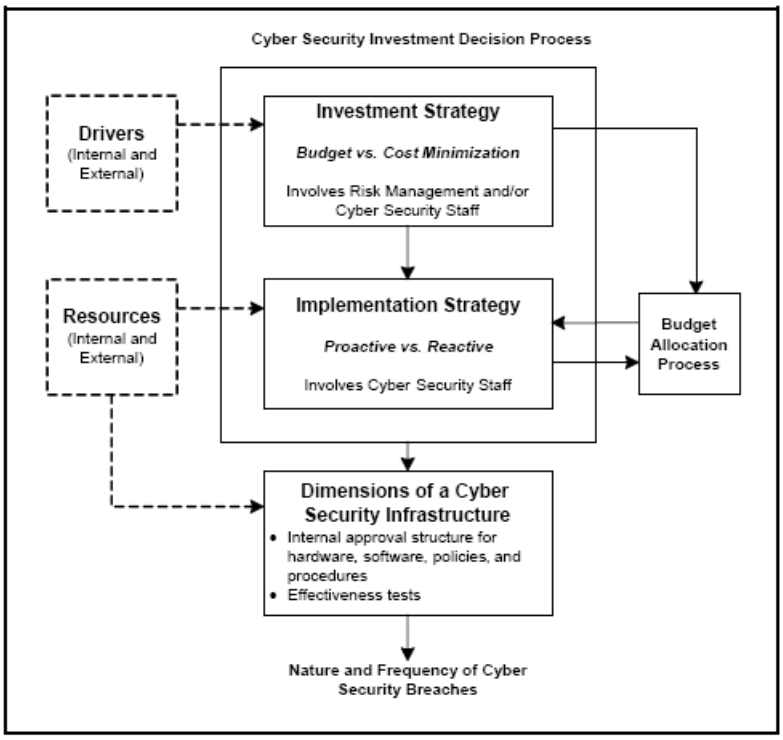

Figure 2.11 Diagram of Cybersecurity Investment by Row and Gallaher [117]