Phishing and Privacy Notices

Why John (Podesta) can’t protect his emails?

Someone has your password Inbox x

Google <no-reply@accounts.googlemail.com> 12:02 (1 hour ago)

to me

Google

Someone has your password

Hi John

Someone just used your password to try to sign in to your Google Account john.podesta@gmail.com.

Details:
Saturday, 19 March, 8:34:30 UTC
IP Address: 134.249.139.239
Location: Ukraine

Google stopped this sign-in attempt. You should change your password immediately.

CHANGE PASSWORD

Bitly Link

THE TRAIL THAT LEADS TO FANCY BEAR

The phishing email that Podesta received on March 19 contained a URL, created with the popular Bitly shortening service, pointing to a longer URL that, to an untrained eye, looked like a Google link.

http://myaccount.google.com-securitysettingpage.tk/

Spam/Phishing/Spear Phishing

- Spam
  - Large scale advertising campaigns via email
  - May want to defraud – but may also want to legitimately sell something
- Phishing
  - Fraudulent email that looks real
  - Usually used to try and extract some information or money
- Spear Phishing
  - A phishing campaign that is targeted and personalized
  - John Podesta Attack

Spam

Where was all this Spam coming from?

- Open Mail Relays
  - Web servers running SMTP that will send/relay any message without authentication
  - Used to be quite common – less so now
- Botnets
  - Networks of infected machines under a centralized control
  - Used for Denial-of-Service attacks
  - Also used to run/send millions of SPAM emails

Who clicks on SPAM?
How can this be profitable?

CCS 2008

Spamalytics: An Empirical Analysis of Spam Marketing Conversion

Chris Kanich* Christian Kreibich† Kirill Levchenko* Brandon Enright*
Geoffrey M. Voelker* Vern Paxson† Stefan Savage*

†International Computer Science Institute
Berkeley, USA
christian@icir.org,vern@cs.berkeley.edu

*Dept. of Computer Science and Engineering
University of California, San Diego, USA
{ckanich,klevchen,voelker,savage}@cs.ucsd.edu
bmenrigh@ucsd.edu

ABSTRACT

The "conversion rate" of spam — the probability that an unsolicited e-mail will ultimately elicit a "sale" — underlies the entire spam value proposition. However, our understanding of this critical behavior is quite limited, and the literature lacks any quantitative study concerning its true value. In this paper we present a methodology for measuring the conversion rate of spam. Using a parasitic infiltration of an existing botnet's infrastructure, we analyze two spam campaigns: one designed to propagate a malware Trojan, the other marketing on-line pharmaceuticals. For nearly a half billion spam e-mails we identify the number that are successfully delivered, the number that pass through popular anti-spam filters, the number that elicit user visits to the advertised sites, and the number of "sales" and "infections" produced.

Unraveling such questions is *essential* for understanding the economic support for spam and hence where any structural weaknesses may lie. Unfortunately, spammers do not file quarterly financial reports, and the underground nature of their activities makes third-party data gathering a challenge at best. Absent an empirical foundation, defenders are often left to speculate as to how successful spam campaigns are and to what degree they are profitable. For example, IBM's Joshua Corman was widely quoted as claiming that spam sent by the Storm worm alone was generating "millions and millions of dollars every day" [2]. While this claim could in fact be true, we are unaware of any public data or methodology capable of confirming or refuting it.

The key problem is our limited visibility into the three basic parameters of the spam value proposition: the cost to send spam, offset by the "conversion rate" (probability that an e-mail sent will

Storm Botnet

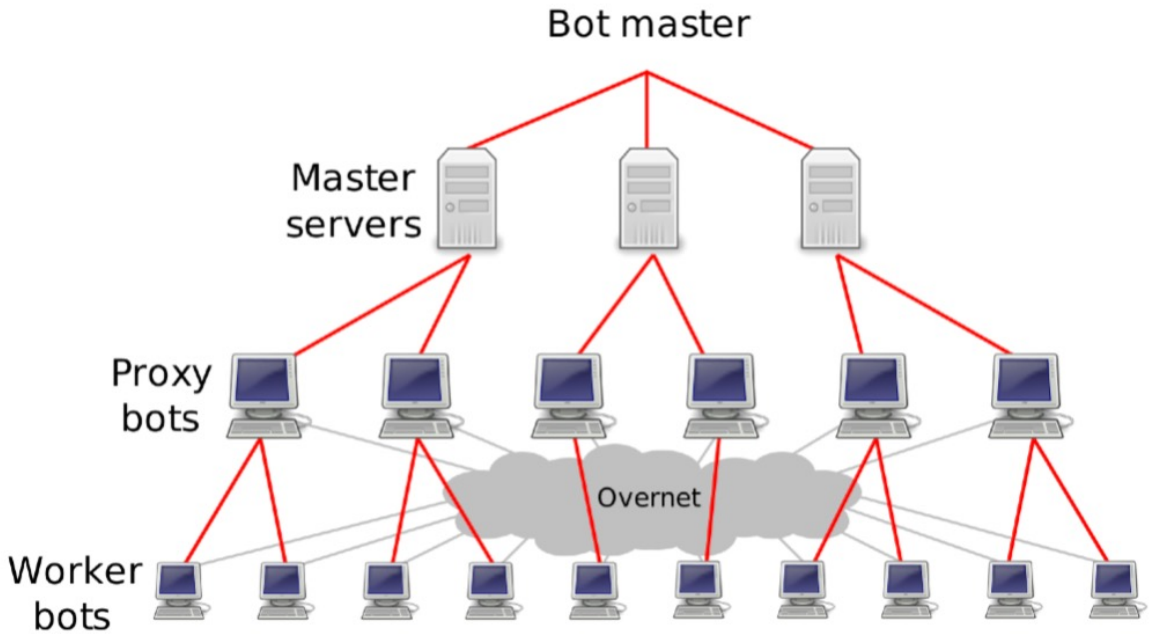

Figure 1: The Storm botnet hierarchy.

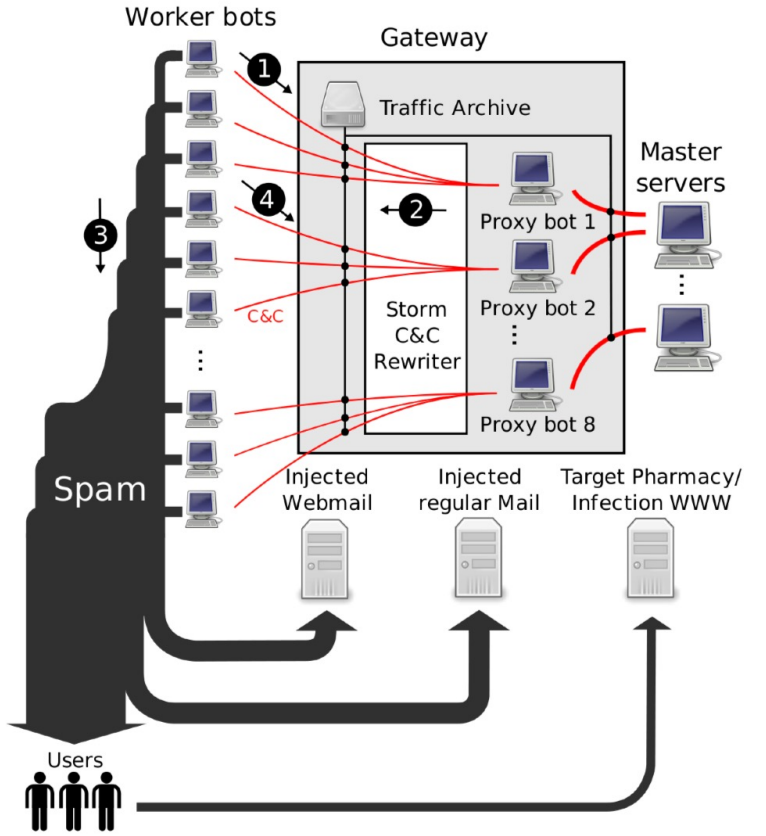

Figure 2: The Storm spam campaign dataflow (Section 3.3) and our measurement and rewriting infrastructure (Section 4). (1) Workers request spam tasks through proxies, (2) proxies forward spam workload responses from master servers, (3) workers send the spam and (4) return delivery reports. Our infrastructure infiltrates the C&C channels between workers and proxies.

What to do when you have a botnet?

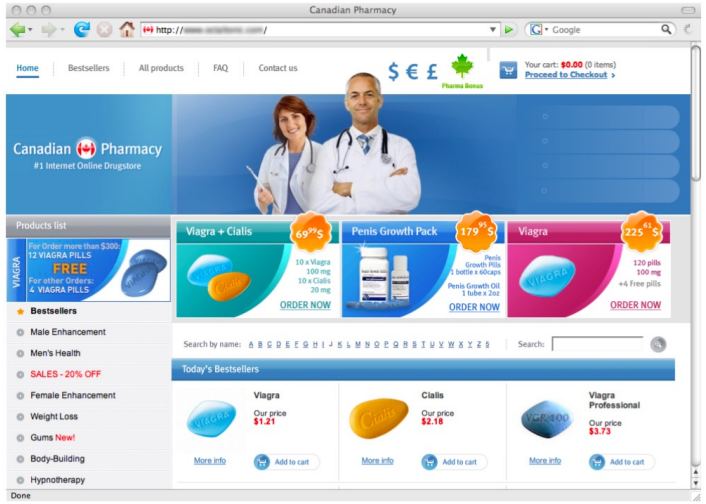

(a) Pharmaceutical site

(b) Postcard-themed self-propagation site

Figure 3: Screenshots of the Web sites operated to measure user click-through and conversion.

4. METHODOLOGY

Our measurement approach is based on *botnet infiltration* — that is, insinuating ourselves into a botnet's "command and control" (C&C) network, passively observing the spam-related commands and data it distributes and, where appropriate, actively changing individual elements of these messages in transit. Storm's architecture lends itself particularly well to infiltration since the proxy bots, *by design*, interpose on the communications between individual worker bots and the master servers who direct them. Moreover, since Storm compromises hosts indiscriminately (normally using malware distributed via social engineering Web sites) it is straightforward to create a proxy bot on demand by infecting a globally reachable host under our control with the Storm malware.

Figure 2 also illustrates our basic measurement infrastructure. At the core, we instantiate eight unmodified Storm proxy bots within a controlled virtual machine environment hosted on VMWare ESX 3 servers. The network traffic for these bots is then routed through a centralized gateway, providing a means for blocking unanticipated behaviors (e.g., participation in DDoS attacks) and an interposition point for parsing C&C messages and "rewriting" them as they pass from proxies to workers. Most critically, by carefully rewriting the spam template and dictionary entries sent by master servers, we arrange for worker bots to replace the intended site links in their spam with URLs of our choosing. From this basic capability we synthesize experiments to measure the click-through and conversion rates for several large spam campaigns.

Conversion Rate

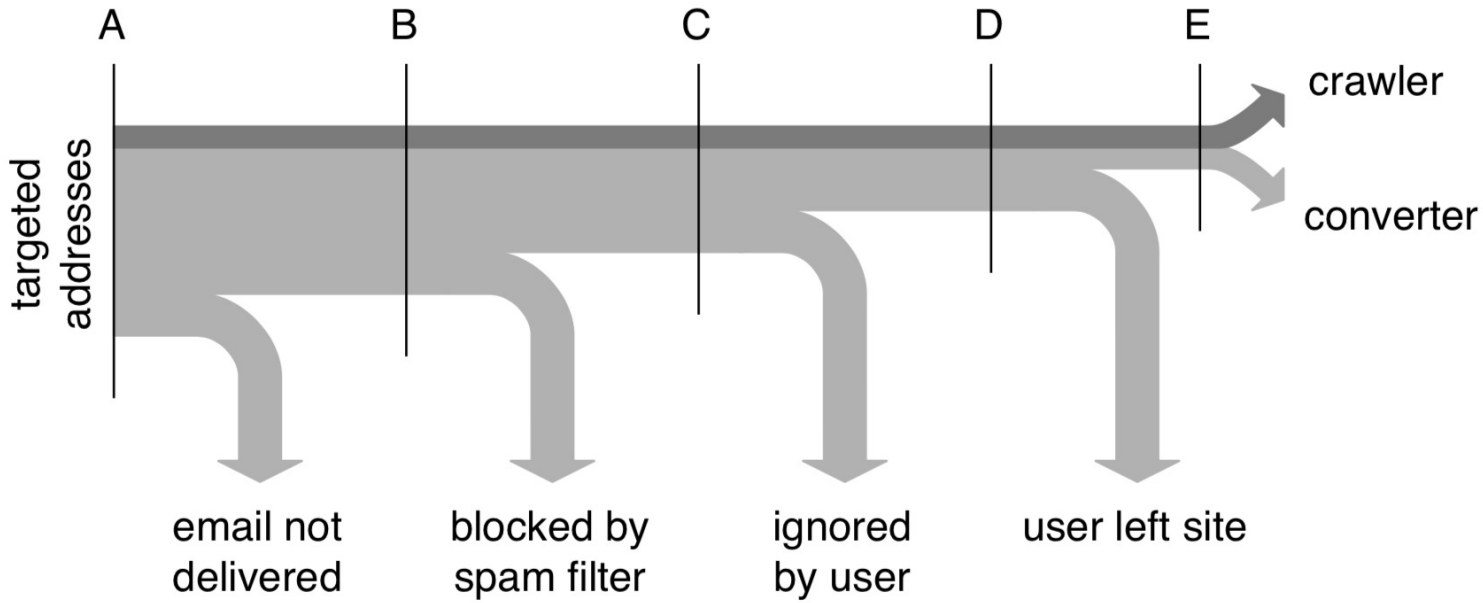

Figure 6: The spam conversion pipeline.

STAGE	PHARMACY		POSTCARD		APRIL FOOL	
A – Spam Targets	347,590,389	100%	83,655,479	100%	40,135,487	100%
B – MTA Delivery (est.)	82,700,000	23.8%	21,100,000	25.2%	10,100,000	25.2%
C – Inbox Delivery	—	—	—	—	—	—
D – User Site Visits	10,522	0.00303%	3,827	0.00457%	2,721	0.00680%
E – User Conversions	28	0.0000081%	316	0.000378%	225	0.000561%

Table 3: Filtering at each stage of the spam conversion pipeline for the self-propagation and pharmacy campaigns. Percentages refer to the conversion rate relative to Stage A.

Profit

After 26 days, and almost 350 million e-mail messages, only 28 sales resulted — a conversion rate of well under 0.00001%. Of these, all but one were for male-enhancement products and the average purchase price was close to $100. Taken together, these conversions would have resulted in revenues of $2,731.88 — a bit over $100 a day for the measurement period or $140 per day for periods when the campaign was active. However, our study interposed on only a small fraction of the overall Storm network — we estimate roughly 1.5 percent based on the fraction of worker bots we proxy. Thus, the total daily revenue attributable to Storm's pharmacy campaign is likely closer to $7000 (or $9500 during periods of campaign activity). By the same logic, we estimate that Storm self-propagation campaigns can produce between 3500 and 8500 new bots per day.

The next obvious question is, "How much of this revenue is profit"? Here things are even murkier. First, we must consider how much of the gross revenue is actually recovered on a sale. Assuming the pharmacy campaign drives traffic to an affiliate program (and there are very strong anecdotal reasons to believe this is so) then the gross revenue is likely split between the affiliate and the program (a annual net revenue of $1.75M using our previous estimate). Next, we must subtract business costs. These include a number of incidental expenses (domain registration, bullet-proof hosting fees, etc) that are basically fixed sunk costs, and the cost to distribute the spam itself.

Spam feels “solved”

- Economics Changed
  - Easier to find online a lot of the items being sold, such as drugs
  - Users “wised up”
- Closing Email Relays
  - Open relays are mostly blocked or tracked
  - More tools to track where emails originated friends
- Natural Language Analysis
  - Bayesian analysis of text got very good at detecting spam
  - Filtering got a lot better

Spam is different than Phishing

- Phishing
  - Designed specifically for deception
  - Sent from real email address, normal language
  - Smaller scale
  - Take advantage of human failures
  - Still a real problem

Phishing

Why does phishing work?

- Lack of Knowledge
  - Lack of computer system knowledge
  - Lack of knowledge of security indicators
- Visual Deceptions
  - Visually deceptive text
  - Images masking underlying text
  - Images mimicking windows
  - Windows making underlying windows
  - Deceptive look and feel
- Bounded Attention
  - Lack of attention to security indicators
  - Lack of attention to the absence of security indicators

CHI 2006

Why Phishing Works

Rachna Dhamija
rachna@deas.harvard.edu
Harvard University

J. D. Tygar
tygar@berkeley.edu
UC Berkeley

Marti Hearst
hearst@sims.berkeley.edu
UC Berkeley

ABSTRACT
To build systems shielding users from fraudulent (or *phishing*) websites, designers need to know which attack strategies work and why. This paper provides the first empirical evidence about which malicious strategies are successful at deceiving general users. We first analyzed a large set of captured phishing attacks and developed a set of hypotheses about why these strategies might work. We then assessed these hypotheses with a usability study in which 22 participants were shown 20 web sites and asked to determine which ones were fraudulent. We found that 23% of the participants did not look at browser-based cues such as the address bar, status bar and the security indicators, leading to incorrect choices 40% of the time. We also found that some visual deception attacks can fool even the most sophisticated users. These results illustrate that standard security indicators are not effective for a substantial fraction of users, and suggest that alternative approaches are needed.

Author Keywords
Security Usability, Phishing.

ACM Classification Keywords
H.1.2 [User/Machine Systems]: Software psychology; K.4.4 [Electronic Commerce]: Security.

INTRODUCTION
What makes a web site credible? This question has been addressed extensively by researchers in computer-human interaction. This paper examines a twist on this question: what makes a *bogus* website credible? In the last two years, Internet users have seen the rapid expansion of a scourge on the Internet: *phishing*, the practice of directing users to fraudulent web sites. This question raises fascinating questions for user interface designers, because both phishers and anti-phishers do battle in user interface space. Successful phishers must not only present a high-credibility web presence to their victims; they must create a presence that is so impressive that it causes the victim to fail to recognize security measures installed in web browsers.

Data suggest that some phishing attacks have convinced up to 5% of their recipients to provide sensitive information to spoofed websites [21]. About two million users gave information to spoofed websites resulting in direct losses of $1.2 billion for U.S. banks and card issuers in 2003 [20].[1]

If we hope to design web browsers, websites, and other tools to shield users from such attacks, we need to understand which attack strategies are successful, and what proportion of users they fool. However, the literature is sparse on this topic.

Visual Indicators (circa 2006)

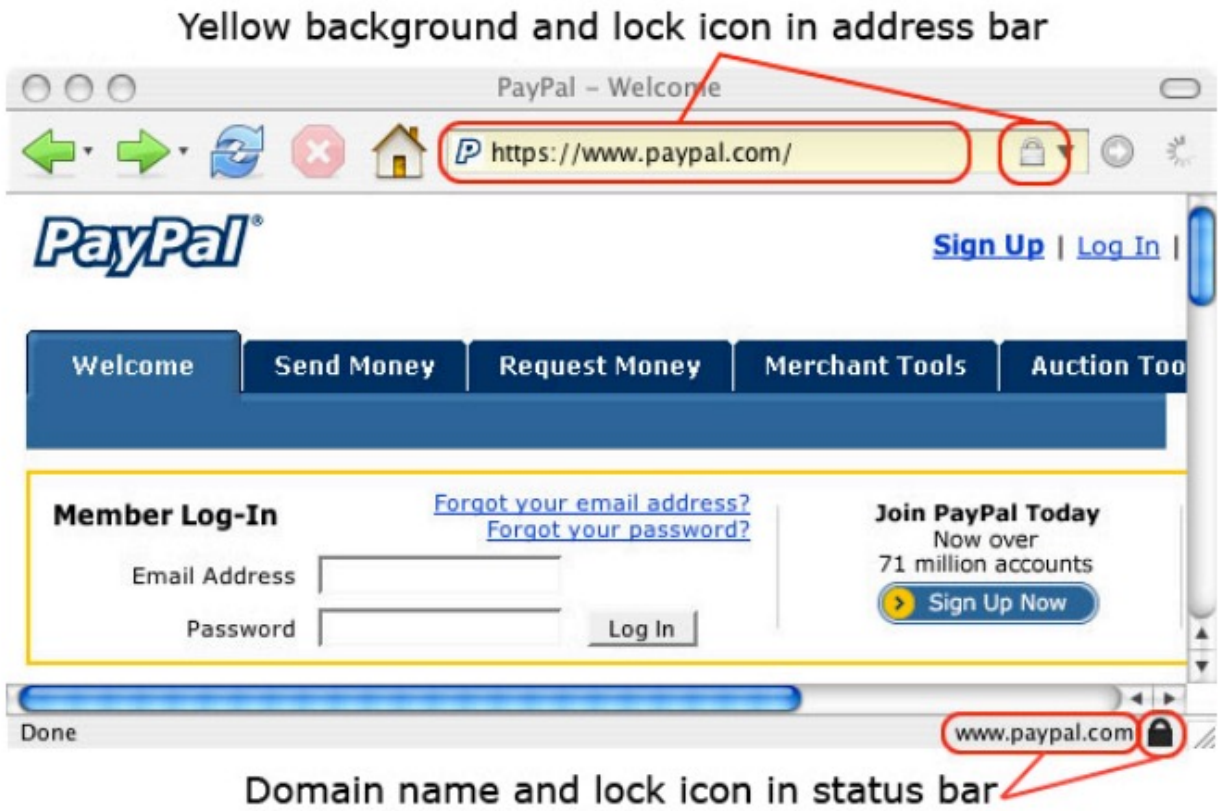

Figure 1: Visual Security Indicators in Mozilla Firefox Browser v1.0.1 for Mac OS X.

Phishing or non Phishing?

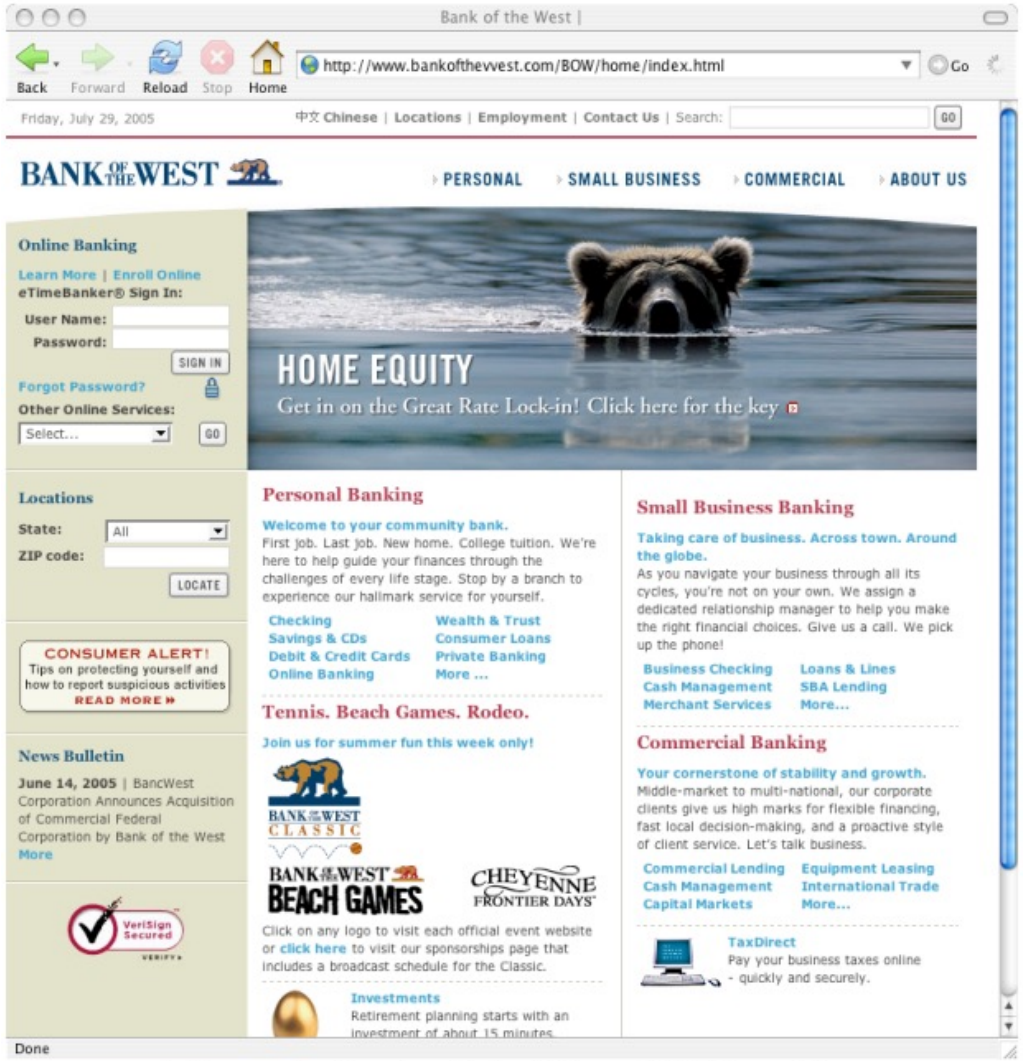

Phishing or non Phishing?

Phishing or non Phishing

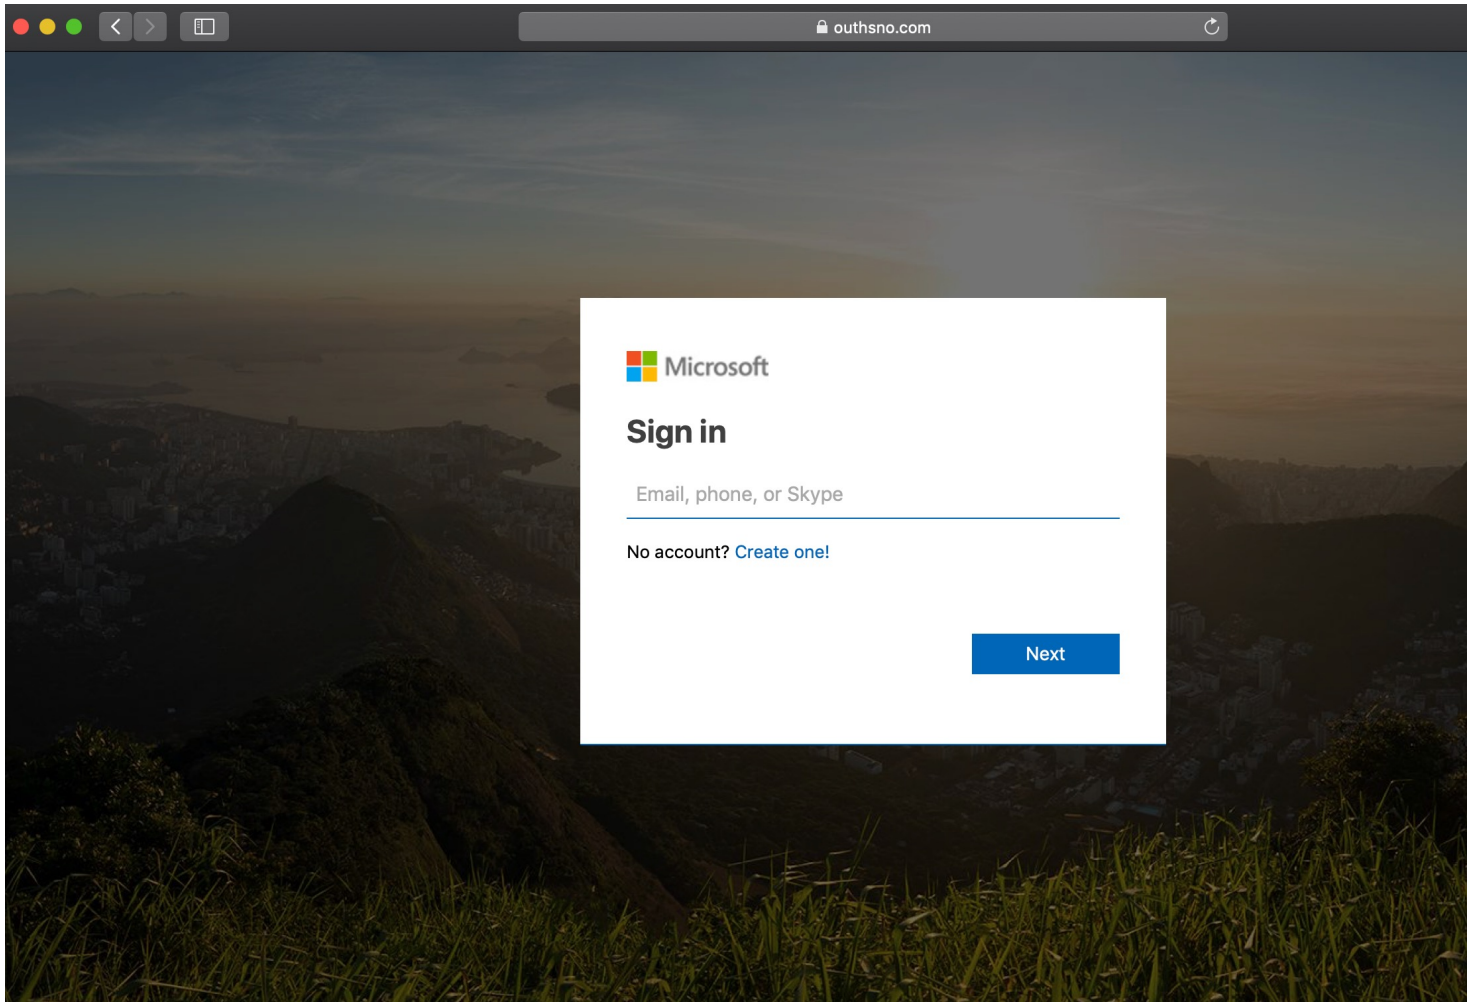

People do bad at these kind of tasks

Website	Real or Spoof	Phishing or Security Tactic Used (Partial List)	% Right (avg conf)	% Wrong (avg conf)
Bank Of the West	Spoof	URL (bankofthevvest.com), padlock in content, Verisign logo and certificate validation seal, consumer alert warning	9 (3.0)	91 (4.2)
PayPal	Spoof	Uses Mozilla XML User Interface Language (XUL) to simulate browser chrome w/ fake address bar, status bar and SSL indicators	18 (3.0)	81 (4.5)
Etrade	Real	3rd party URL (etrade.everypath.com), SSL, simple design, no graphics for mobile users	23 (4.6)	77 (4.2)
PayPal	Spoof	URL (paypal-signin03.com), padlock in content	41 (4.0)	59 (3.7)
PayPal	Spoof	URL (IP address), padlock in content	41 (3.9)	59 (4.5)
Capital One	Real	3rd party URL (cib.ibanking-services.com), SSL, dedicated login page, simple design	50 (3.9)	50 (3.5)
Paypal	Spoof	Screenshot of legitimate SSL protected Paypal page within a rogue webpage	50 (4.7)	50 (4.3)
Ameritrade	Spoof	URL (ameritrading.net)	50 (4.2)	50 (3.9)
Bank of America	Spoof	Rogue popup window on top of legitimate BOFA homepage, padlock in content	64 (4.2)	36 (4.4)
Bank Of The West	Spoof	URL (IP address), urgent anti-fraud warnings (requests large amount of personal data)	68 (4.8)	32 (4.4)
USBank	Spoof	URL (IP address), padlock in content, security warnings, identity verification (requests large amount of personal data)	68 (4.1)	32 (4.3)
Ebay	Spoof	URL (IP address), account verification (requests large amount of personal data)	68 (4.4)	32 (4.0)
Yahoo	Spoof	URL (center.yahoo-security.net), account verification (requests large amount of personal data)	77 (3.0)	23 (4.2)
NCUA	Spoof	URL (IP address), padlock in content, account verification (requests large amount of personal data)	82 (4.5)	18 (4.3)
Ebay	Real	SSL protected login page, TRUSTe logo	86 (4.4)	14 (4.0)
Bank Of America	Real	Login page on non-SSL homepage, padlock in content	86 (4.4)	14 (3.3)
Tele-Bears (Student Accounts)	Real	SSL protected login page	91 (4.7)	9 (4.5)
PayPal	Real	Login page on non-SSL homepage, padlock in content	91 (4.6)	9 (3.0)
Bank One	Real	Login page on non-SSL homepage, padlock in content	100 (4.0)	0 (N/A)

Table 2: Security or spoofing strategy employed by each site (spoof sites shown with white background, real sites gray).

Phishing is still a big problem today …

PhishTank

https://phishtank.org/

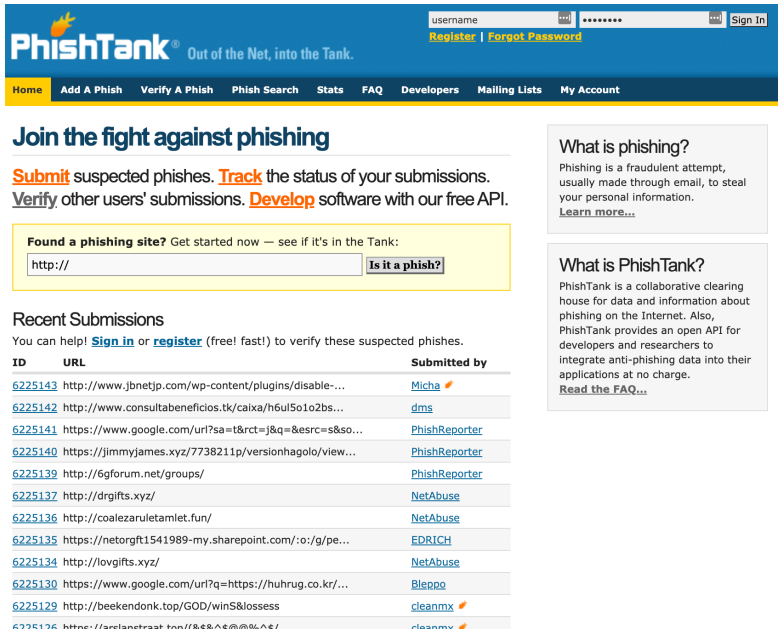

What is PhishTank?

PhishTank is a free community site where anyone can submit, verify, track and share phishing data.

Does PhishTank cost anything?

PhishTank is free to everyone, both the website and the data (via the API).

Does PhishTank protect me from phishing?

PhishTank is not protection. PhishTank is an information clearinghouse, which helps to pour sunshine on some of the dark alleys of the Internet. PhishTank provides accurate, actionable information to anyone trying to identify bad actors, whether for themselves or for others (i.e., building security tools).

Who is behind PhishTank?

PhishTank is operated by OpenDNS, a company founded in 2005 to improve the Internet through safer, faster, and smarter DNS. Read more at www.opendns.com.

Why does OpenDNS operate PhishTank?

OpenDNS is interested in having the best available information about phishing websites. However, phishing data is not a place to be competitive, and we believe that sharing this data freely (even with those who do not contribute) will benefit us all. PhishTank's mission is in line with both OpenDNS's business and its goal of making the Internet a better place.

Why do I have to register to report a suspected phish?

Registration helps make the data better. PhishTank needs to attribute reporting and validation to individual accounts, so the community can learn to judge each member's contribution. This small hurdle also reduces "noise" in the submissions. You are not asked for a lot of personal information: a valid email address is the only personally-identifiable information required. PhishTank needs to attribute reporting and validation to individual accounts, so the community can learn to judge each member's contribution.

How do I report a suspected phish via email?

Submissions via email are strongly encouraged, as more data is usually available. After completing the free registration, you can send emails to phish@phishtank.com from your registered email address. It is important to include as much information as possible, including mail headers if possible. For that reason, we suggest **redirecting** any suspected phishes to PhishTank. To submit suspected phishes from other email addresses, use your individual phish reporting address, which is available from My Account page once you are signed in. We suggest adding your individual phish reporting address to your address book in every mail application you use, for all accounts.

PhishTank

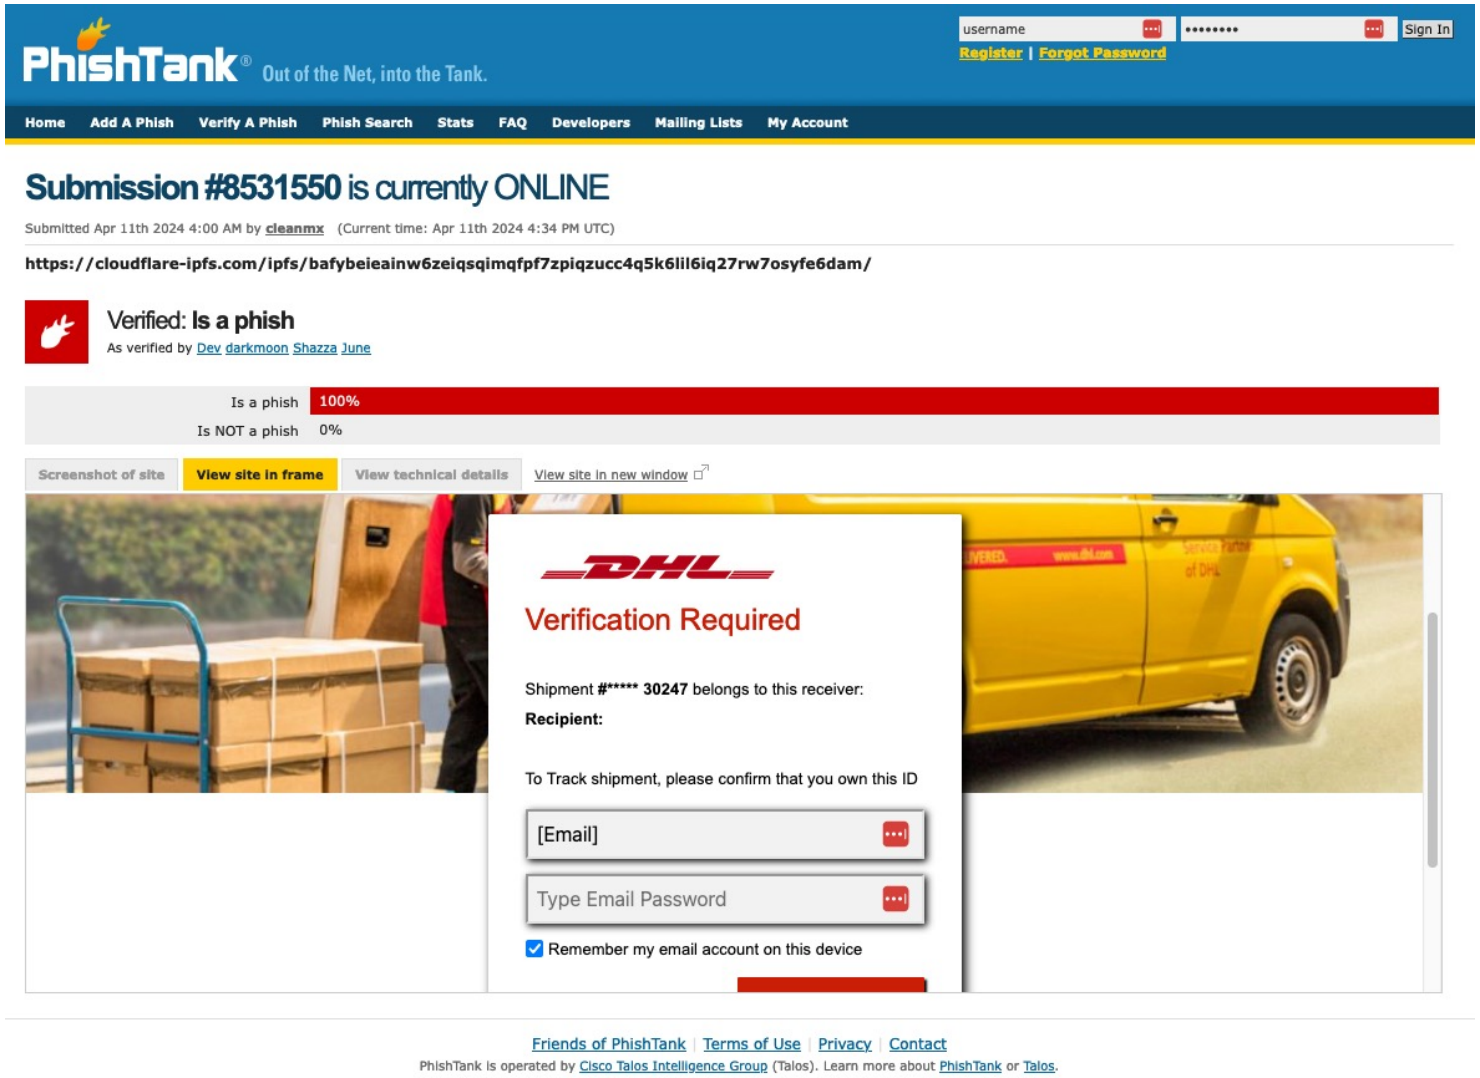
PhishTank
Out of the Net, into the Tank.
username
Sign In
Register | Forgot Password
Home
Add A Phish
Verify A Phish
Phish Search
Stats
FAQ
Developers
Mailing Lists
My Account
Submission #8531550 is currently ONLINE
Submitted Apr 11th 2024 4:00 AM by cleanmx (Current time: Apr 11th 2024 4:34 PM UTC)
https://cloudflare-ipfs.com/ipfs/bafybeieainw6zeiqsqimqfpf7zpiqzucc4q5k6lil6iq27rw7osyfe6dam/
Verified: Is a phish
As verified by Dev darkmoon Shazza June
Is a phish 100%
Is NOT a phish 0%
Screenshot of site
View site in frame
View technical details
View site in new window
DHL
Verification Required
Shipment #***** 30247 belongs to this receiver:
Recipient:
To Track shipment, please confirm that you own this ID
[Email]
Type Email Password
Remember my email account on this device
Friends of PhishTank | Terms of Use | Privacy | Contact
PhishTank is operated by Cisco Talos Intelligence Group (Talos). Learn more about PhishTank or Talos.

PhishTank

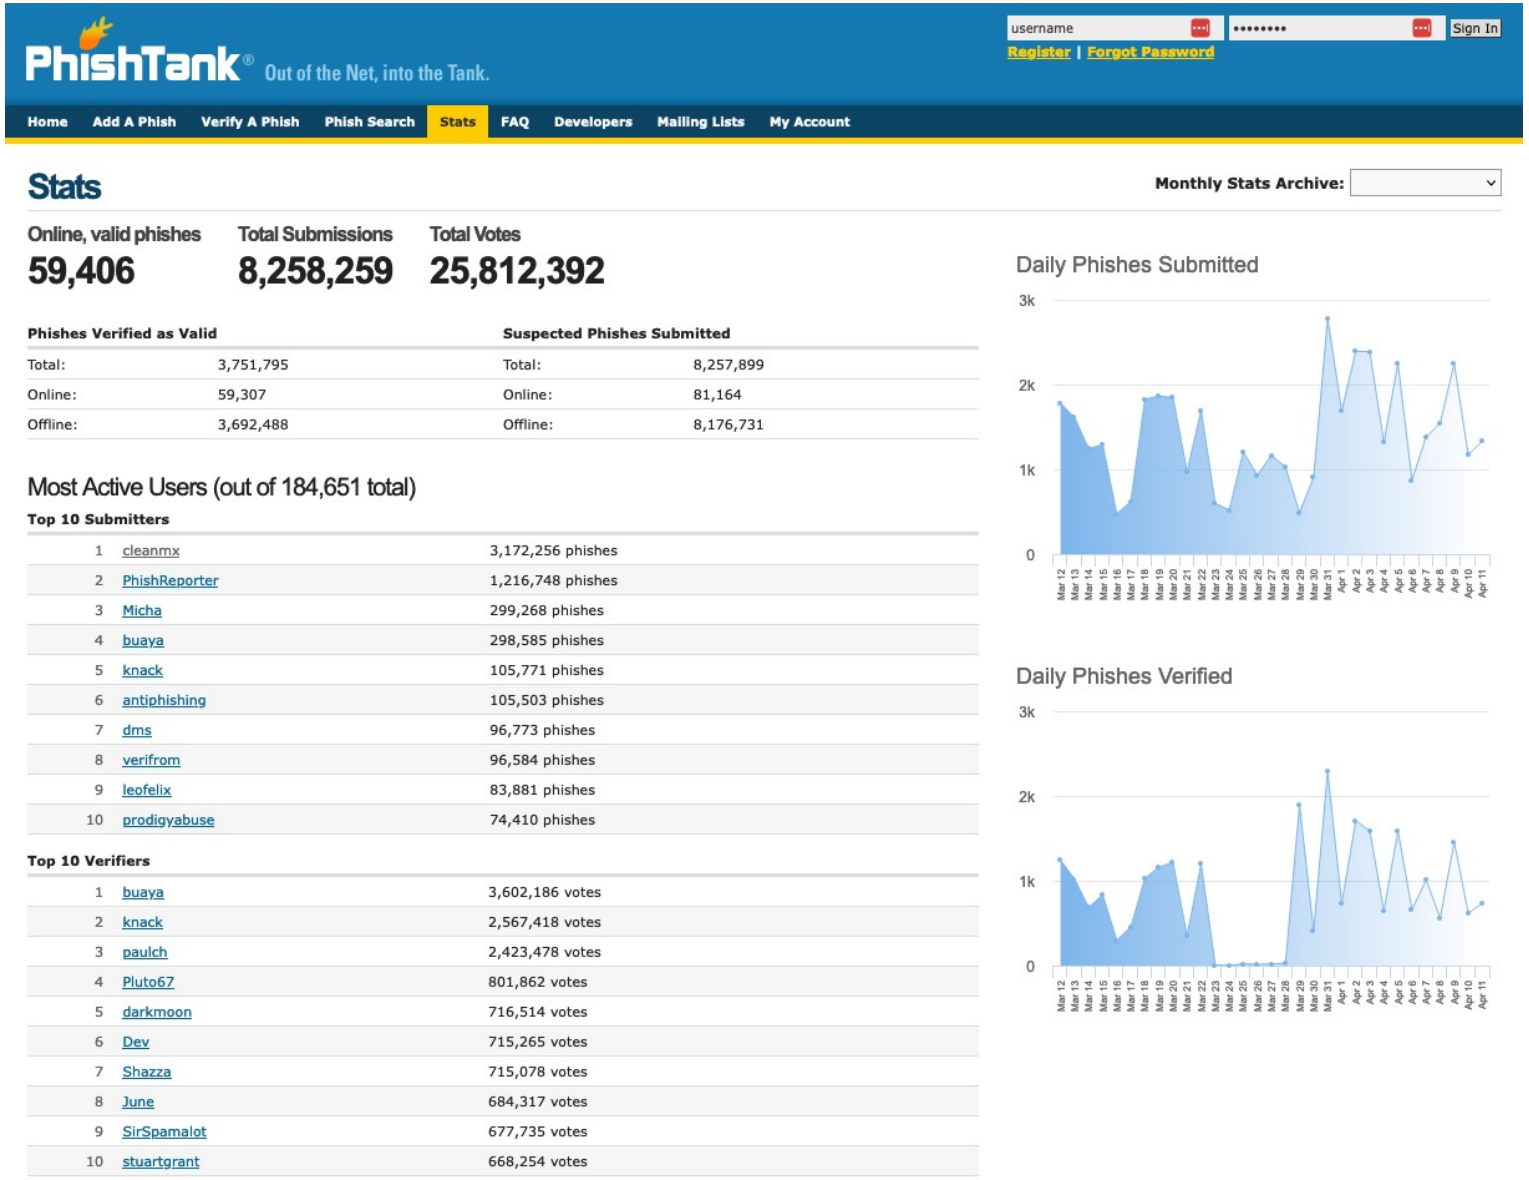
PhishTank® Out of the Net, into the Tank.
username
Sign In
Register | Forgot Password
Home
Add A Phish
Verify A Phish
Phish Search
Stats
FAQ
Developers
Mailing Lists
My Account
Stats
Monthly Stats Archive:
Online, valid phishes
59,406
Total Submissions
8,258,259
Total Votes
25,812,392
Phishes Verified as Valid
Total: 3,751,795
Online: 59,307
Offline: 3,692,488
Suspected Phishes Submitted
Total: 8,257,899
Online: 81,164
Offline: 8,176,731
Most Active Users (out of 184,651 total)
Top 10 Submitters
1 cleanmx 3,172,256 phishes
2 PhishReporter 1,216,748 phishes
3 Micha 299,268 phishes
4 buaya 298,585 phishes
5 knack 105,771 phishes
6 antiphishing 105,503 phishes
7 dms 96,773 phishes
8 verifrom 96,584 phishes
9 leofelix 83,881 phishes
10 prodigyabuse 74,410 phishes
Top 10 Verifiers
1 buaya 3,602,186 votes
2 knack 2,567,418 votes
3 paulch 2,423,478 votes
4 Pluto67 801,862 votes
5 darkmoon 716,514 votes
6 Dev 715,265 votes
7 Shazza 715,078 votes
8 June 684,317 votes
9 SirSpamalot 677,735 votes
10 stuartgrant 668,254 votes
Daily Phishes Submitted
3k
2k
1k
0
Daily Phishes Verified
3k
2k
1k
0

PhishTank – user verified phishing

Examples: Phishing email / Phishing website

Show email annotations

From: First Generic Bank <accounts@firstgenericbank.com>
Subject: Please update your account information
Date: Sep 12, 2006 3:23 PM PST

Dear First Generic Bank user,

As a courtesy to our valued customers, First Generic Bank conducts regular account information verification processes. During the most recent process, we found that we could not verify your information.

In order to ensure your account information is not made vulnerable, please visit http://www.firstgenericbank.com.account-updateinfo.com.

Please click on the above link to our Web site and confirm or update your account information. If you do not do this within 48 hours of receipt of this e-mail, you will not be able to use your First Generic Bank account for 30 days. This is an extra precaution we take to ensure your account remains secure.

Sincerely,

First Generic Bank

What to look for in a phishing email

1. **Generic greeting.** Phishing emails are usually sent in large batches. To save time, Internet criminals use generic names like "First Generic Bank Customer" so they don't have to type all recipients' names out and send emails one-by-one. If you don't see your name, be suspicious.
2. **Forged link.** Even if a link has a name you recognize somewhere in it, it doesn't mean it links to the real organization. Roll your mouse over the link and see if it matches what appears in the email. If there is a discrepency, don't click on the link. Also, websites where it is safe to enter personal information begin with "https" — the "s" stands for secure. If you don't see "https" do not proceed.
3. **Requests personal information.** The point of sending phishing email is to trick you into providing your personal information. If you receive an email requesting your personal information, it is probably a phishing attempt.
4. **Sense of urgency.** Internet criminals want you to provide your personal information now. They do this by making you think something has happened that requires you to act fast. The faster they get your information, the faster they can move on to another victim.

PhishTank – user verified phishing

Examples: **Phishing email** / Phishing website

Hide email annotations

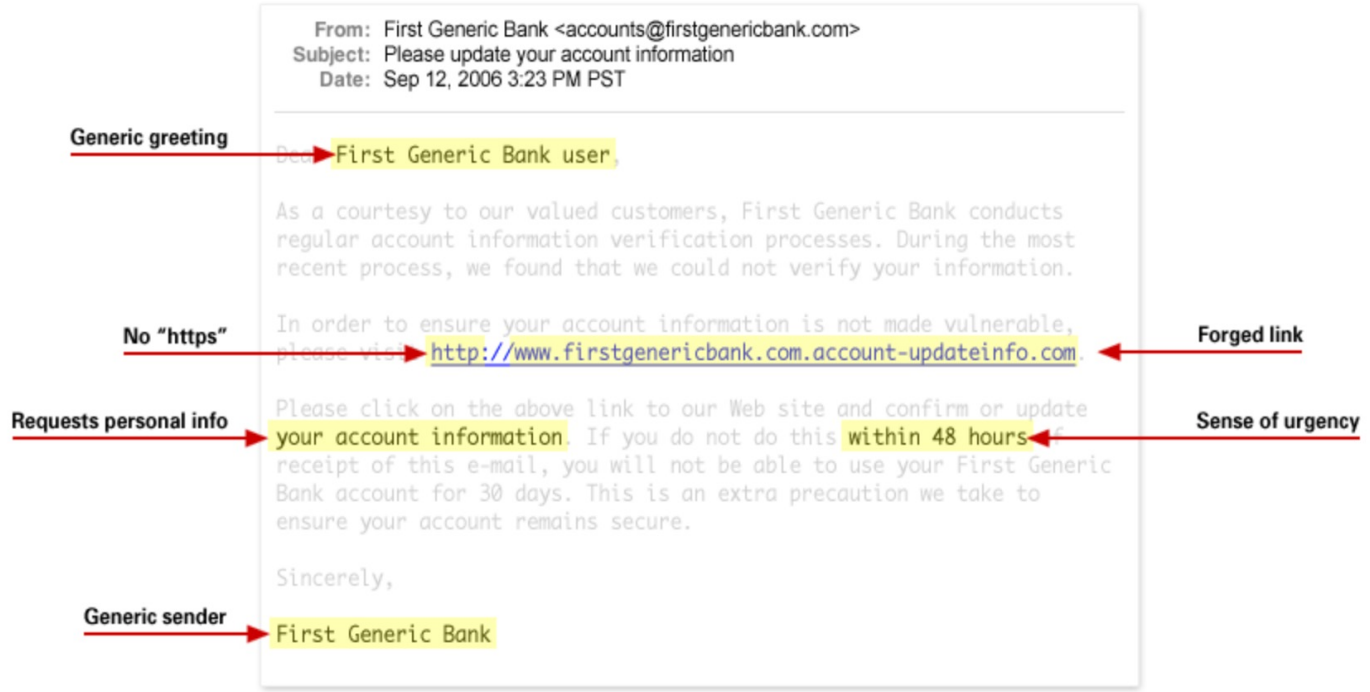

What to look for in a phishing email

1. **Generic greeting.** Phishing emails are usually sent in large batches. To save time, Internet criminals use generic names like "First Generic Bank Customer" so they don't have to type all recipients' names out and send emails one-by-one. If you don't see your name, be suspicious.
2. **Forged link.** Even if a link has a name you recognize somewhere in it, it doesn't mean it links to the real organization. Roll your mouse over the link and see if it matches what appears in the email. If there is a discrepency, don't click on the link. Also, websites where it is safe to enter personal information begin with "https" — the "s" stands for secure. If you don't see "https" do not proceed.
3. **Requests personal information.** The point of sending phishing email is to trick you into providing your personal information. If you receive an email requesting your personal information, it is probably a phishing attempt.
4. **Sense of urgency.** Internet criminals want you to provide your personal information now. They do this by making you think something has happened that requires you to act fast. The faster they get your information, the faster they can move on to another victim.

PhishTank – user verified website

Examples: Phishing email / Phishing website

Show website annotations

What to look for in a phishing website

1. **Poor resolution.** Phishing websites are often poor in quality, since they are created with urgency and have a short lifespan. If the resolution on a logo or in text strikes you as poor, be suspicious.
2. **Forged URL.** Even if a link has a name you recognize somewhere in it, it doesn't mean it links to the real organization. Read URLs from right to left — the real domain is at the end of the URL. Also, websites where it is safe to enter personal information begin with "https" — the "s" stands for secure. If you don't see "https" do not proceed. Look out for URLs that begin with an IP address, such as: http://12.34.56.78/firstgenericbank/account-update/ — these are likely phishes.

PhishTank – user verified website

Examples: Phishing email / Phishing website

Hide website annotations

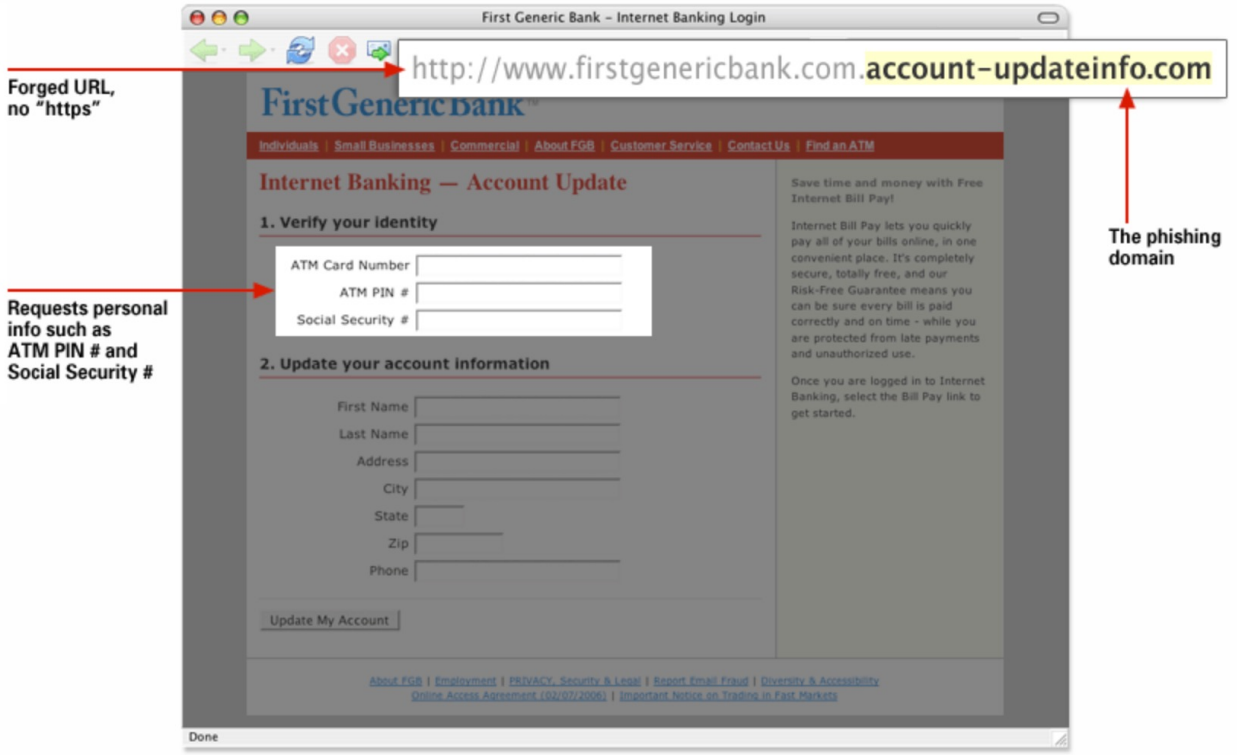

What to look for in a phishing website

1. **Poor resolution.** Phishing websites are often poor in quality, since they are created with urgency and have a short lifespan. If the resolution on a logo or in text strikes you as poor, be suspicious.
2. **Forged URL.** Even if a link has a name you recognize somewhere in it, it doesn't mean it links to the real organization. Read URLs from right to left — the real domain is at the end of the URL. Also, websites where it is safe to enter personal information begin with "https" — the "s" stands for secure. If you don't see "https" do not proceed. Look out for URLs that begin with an IP address, such as: http://12.34.56.78/firstgenericbank/account-update/ — these are likely phishes.

Helping end users?

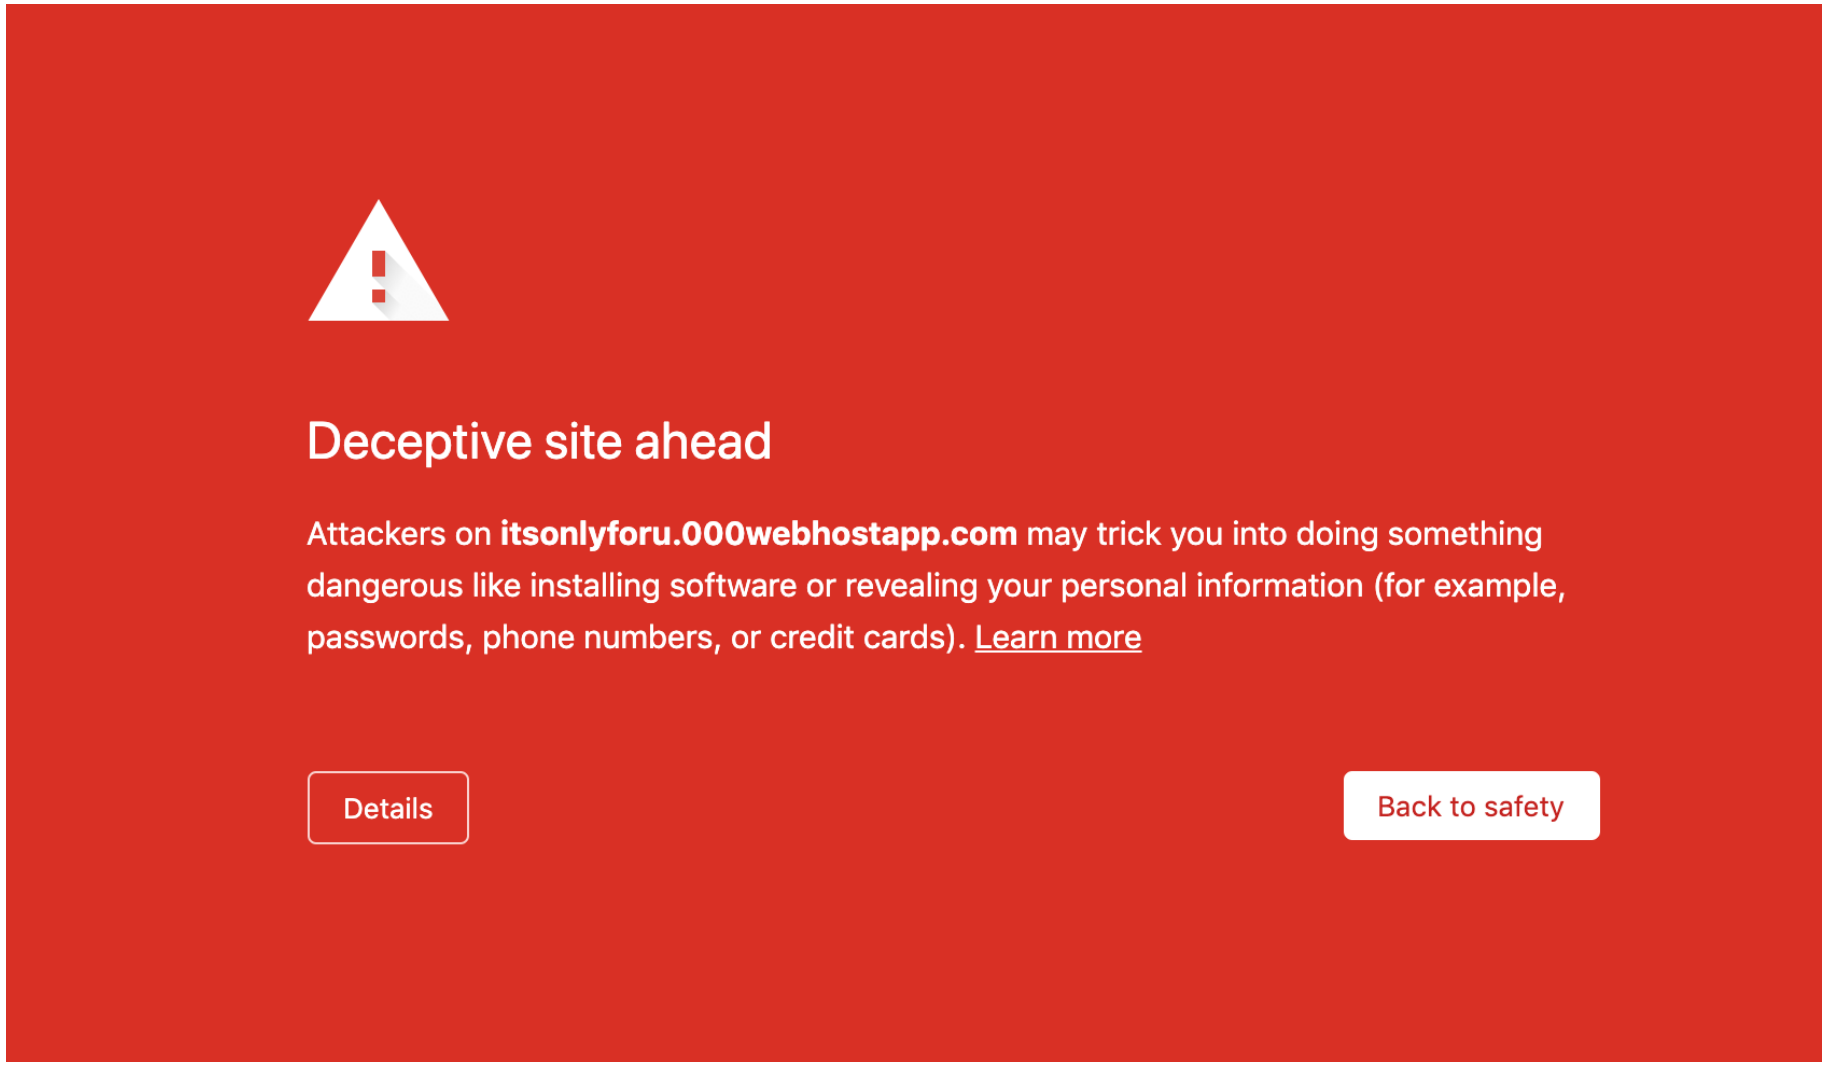

PhishGuru -- Training

PhishGuru: A System for Educating Users about Semantic Attacks

Ponnurangam Kumaraguru

CMU-ISR-09-106
April 14, 2009

Training

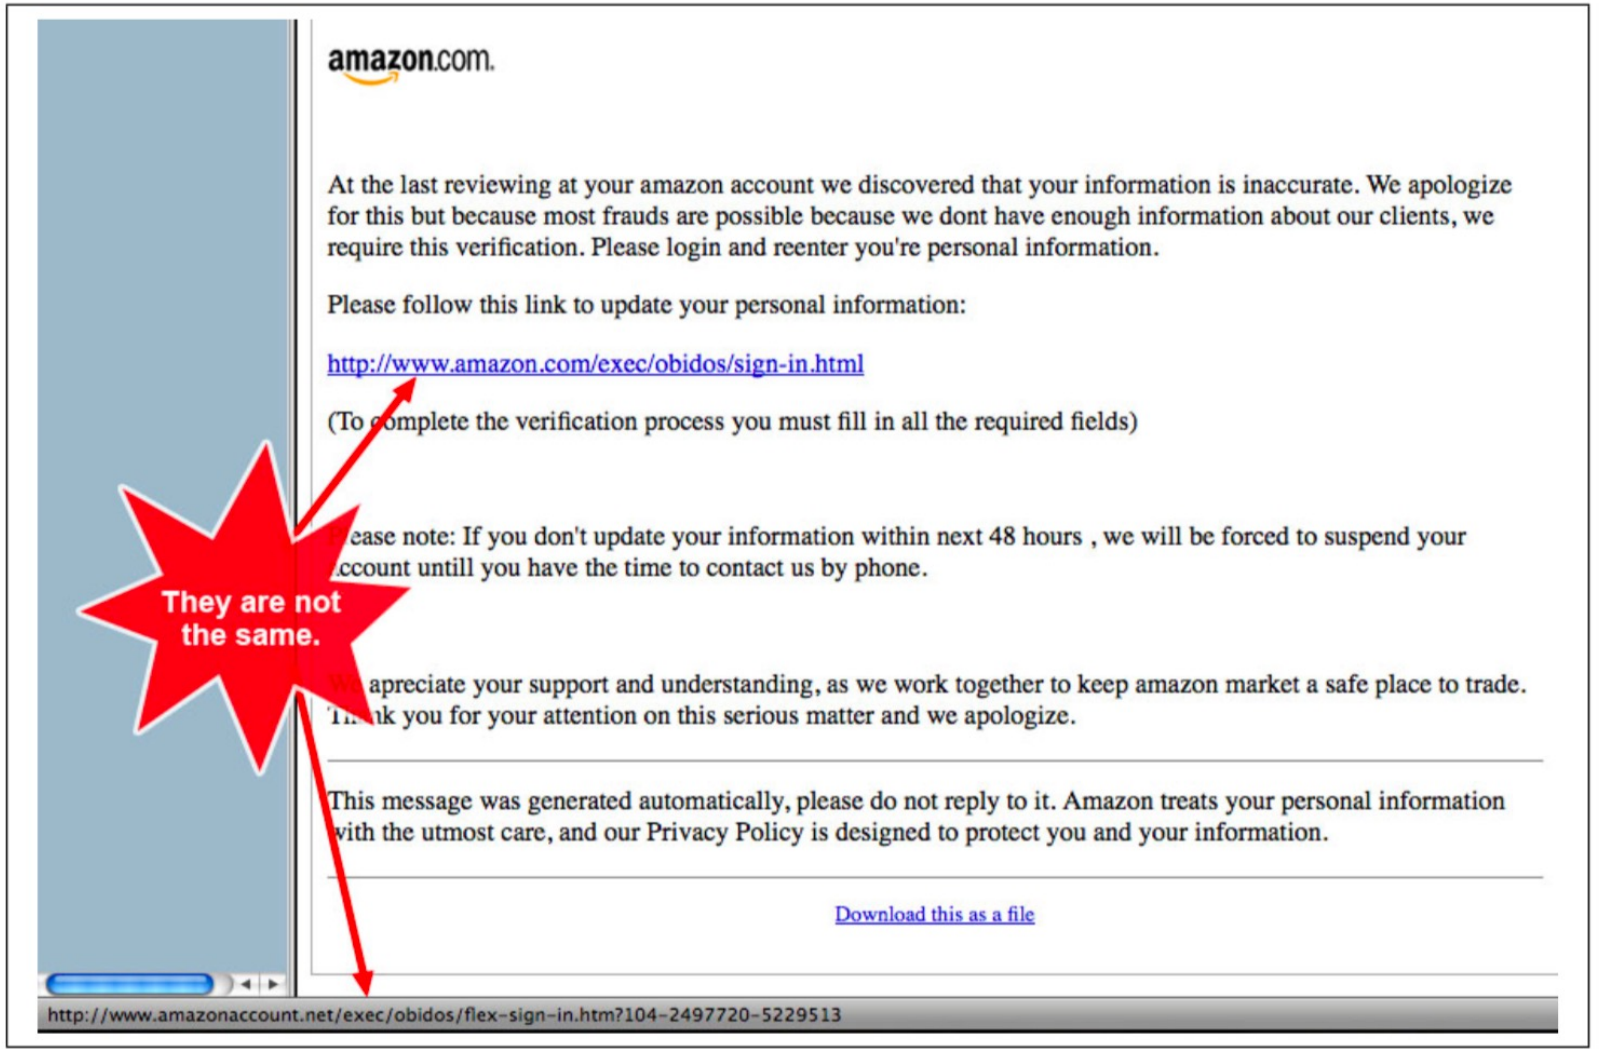
amazon.com.
At the last reviewing at your amazon account we discovered that your information is inaccurate. We apologize for this but because most frauds are possible because we dont have enough information about our clients, we require this verification. Please login and reenter you're personal information.
Please follow this link to update your personal information:
http://www.amazon.com/exec/obidos/sign-in.html
(To complete the verification process you must fill in all the required fields)
ease note: If you don't update your information within next 48 hours , we will be forced to suspend your ccount untill you have the time to contact us by phone.
They are not the same.
preciate your support and understanding, as we work together to keep amazon market a safe place to trade. nk you for your attention on this serious matter and we apologize.
This message was generated automatically, please do not reply to it. Amazon treats your personal information with the utmost care, and our Privacy Policy is designed to protect you and your information.
Download this as a file
http://www.amazonaccount.net/exec/obidos/flex-sign-in.htm?104-2497720-5229513

Training

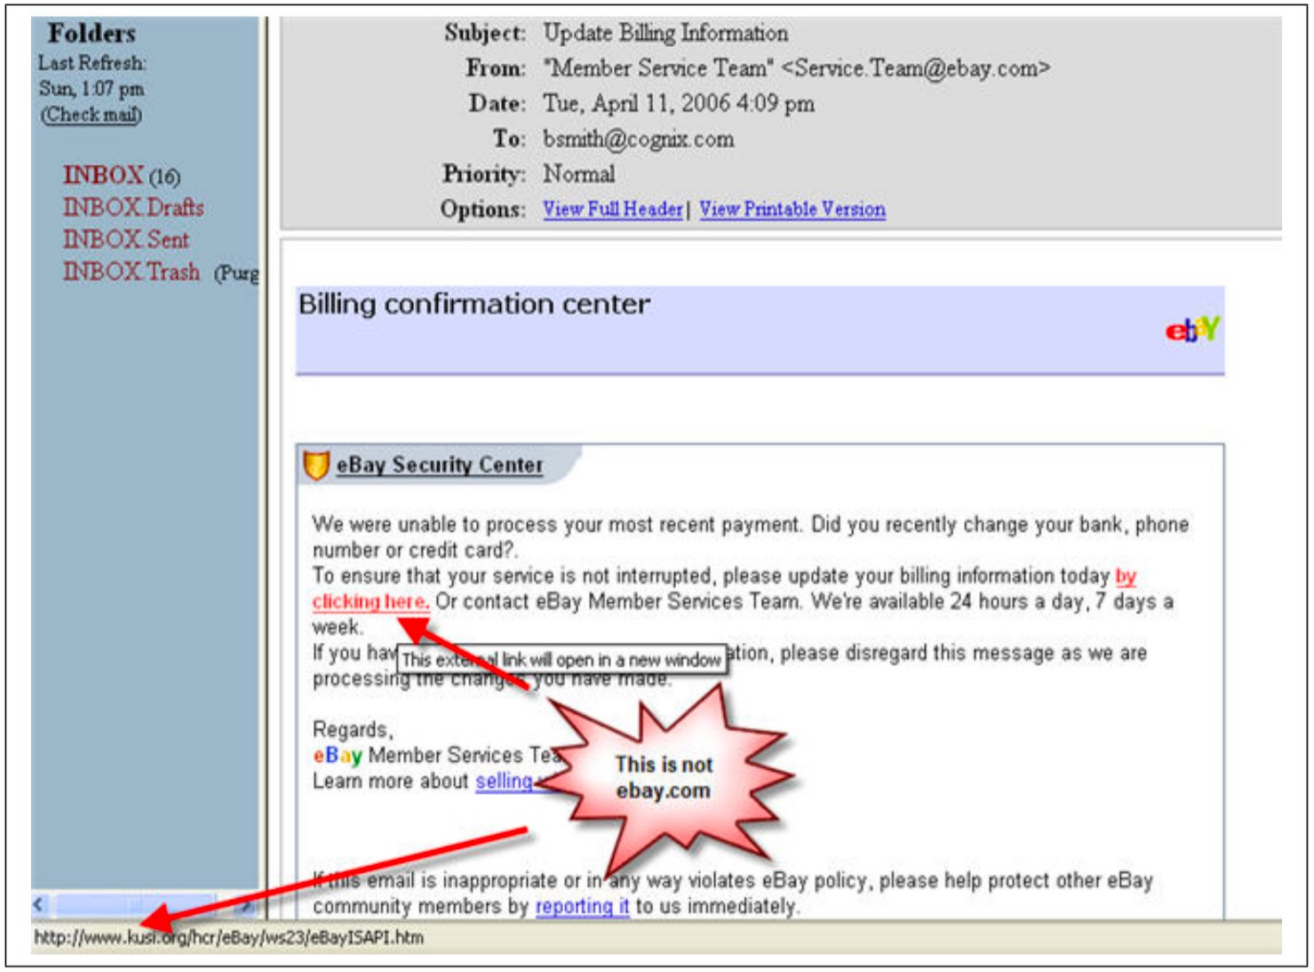
Folders
Last Refresh:
Sun, 1:07 pm
(Check mail)
INBOX (16)
INBOX.Drafts
INBOX.Sent
INBOX.Trash
Subject: Update Billing Information
From: "Member Service Team" <Service.Team@ebay.com>
Date: Tue, April 11, 2006 4:09 pm
To: bsmith@cognix.com
Priority: Normal
Options: View Full Header | View Printable Version
Billing confirmation center
eBay Security Center
We were unable to process your most recent payment. Did you recently change your bank, phone number or credit card?.
To ensure that your service is not interrupted, please update your billing information today by clicking here. Or contact eBay Member Services Team. We're available 24 hours a day, 7 days a week.
This external link will open in a new window
processing the changes you have made.
Regards,
eBay Member Services Te
Learn more about selling
This is not ebay.com
If this email is inappropriate or in any way violates eBay policy, please help protect other eBay community members by reporting it to us immediately.
http://www.kusi.org/hcr/eBay/ws23/eBayISAPI.htm

Training

Protect yourself from Phishing Scams

Clicking on links within emails like the one in the "amazon.com" email you've just read puts you at risk for **identity theft** and **financial loss**.
This email and tutorial were developed by **Carnegie Mellon University** to teach you how to **protect yourself** from these kind of **phishing scams**.

1. What's a phishing scam?

- Scammers send fake emails impersonating well-known companies to trick you into giving them your personal information.
- Giving up your personal information such as Social Security Number, credit card number, or account password will lead to identity theft and financial loss.

2. What does a phishing scam look like?

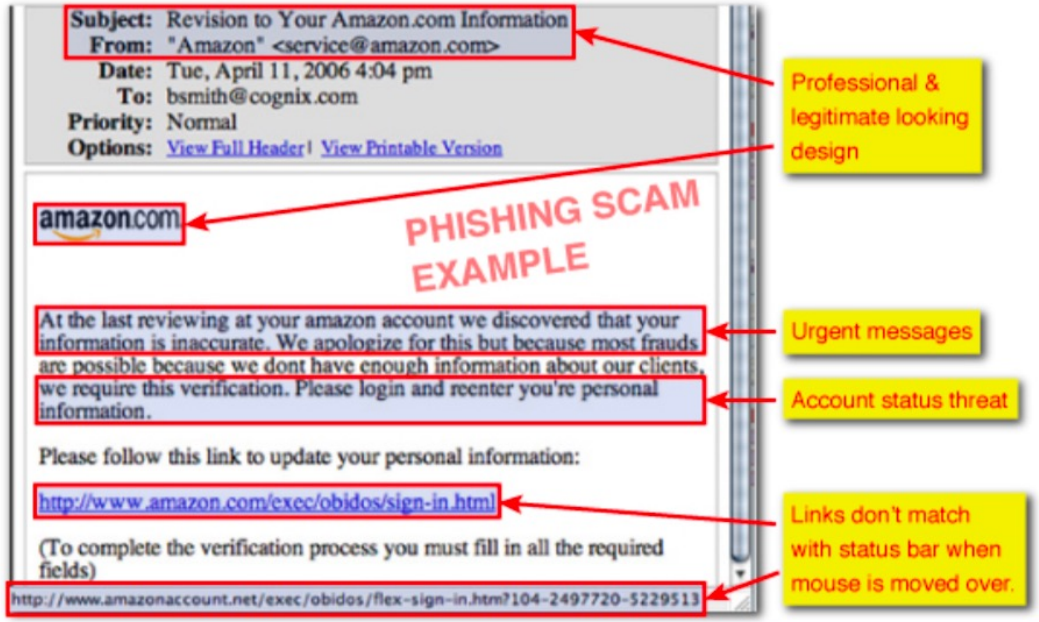

3. What are simple ways to protect yourself from phishing scams?

- **Never click on links within emails:** Never click on links within emails or reply to emails asking for your personal information.
- **Initiate contact:** Always access a website by typing in the real website address into the web browser.

- **Call customer service:** Never trust phone numbers within emails. Look it up yourself and call the customer service when email seems suspicious.
- **Never give out personal information:** Never give out personal information upon email request. Companies will rarely ask for your personal information via emails.

Training (iteration)

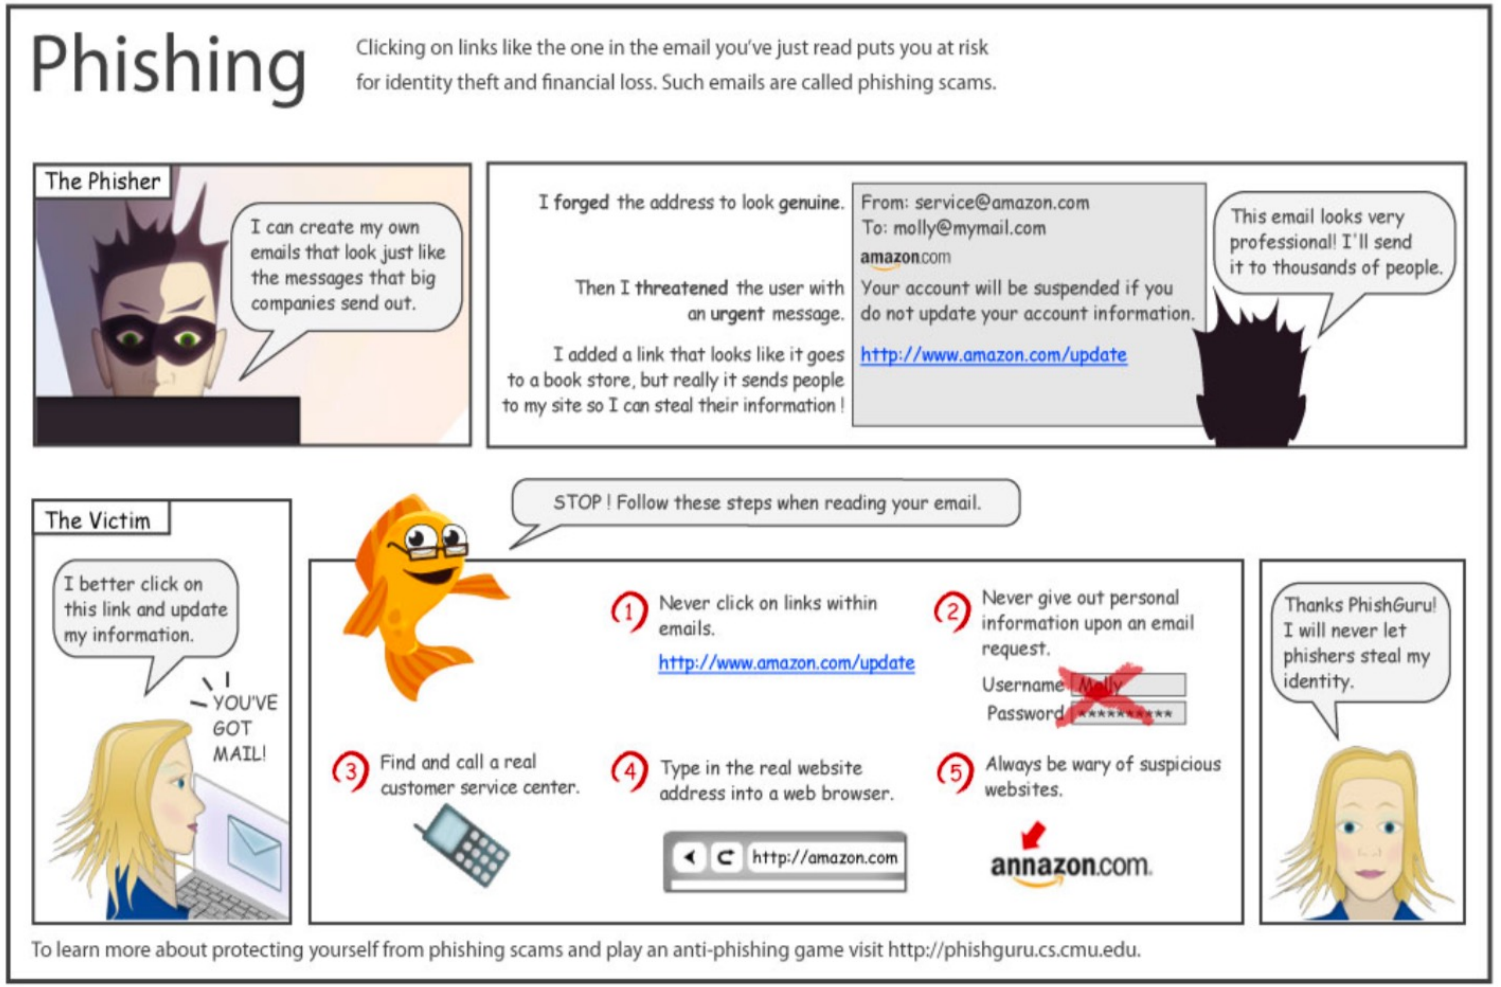

Training (iteration)

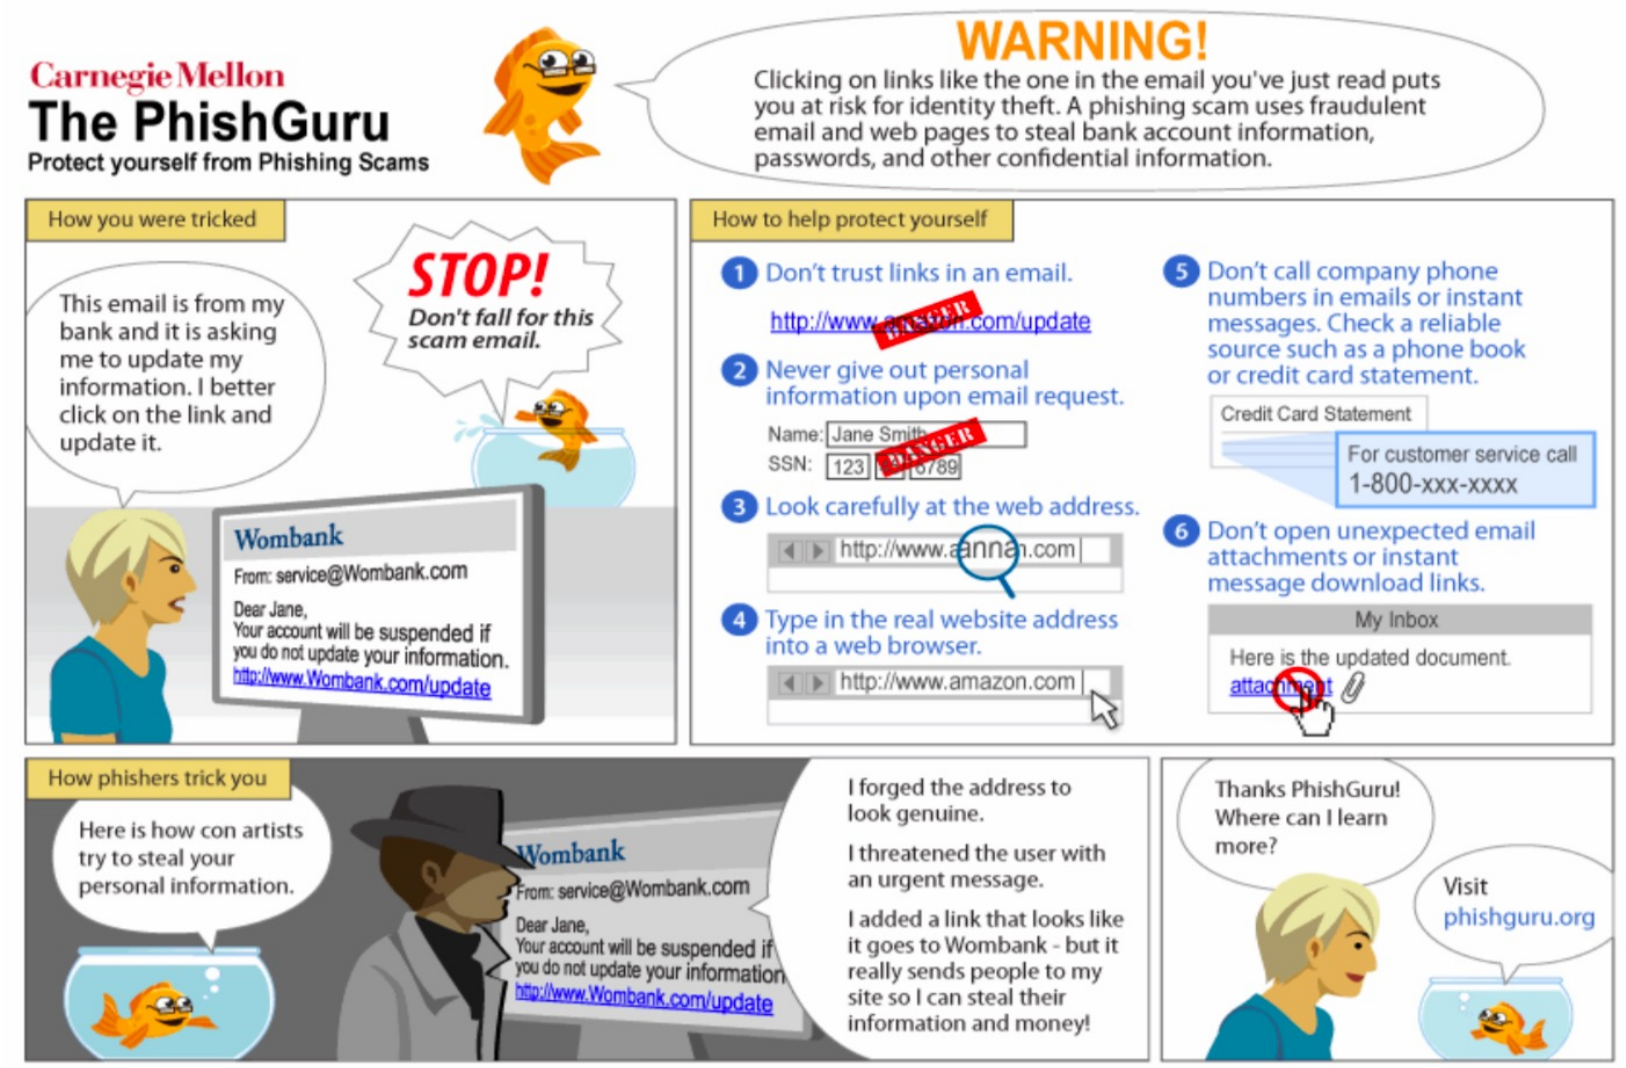

Training (iteration)

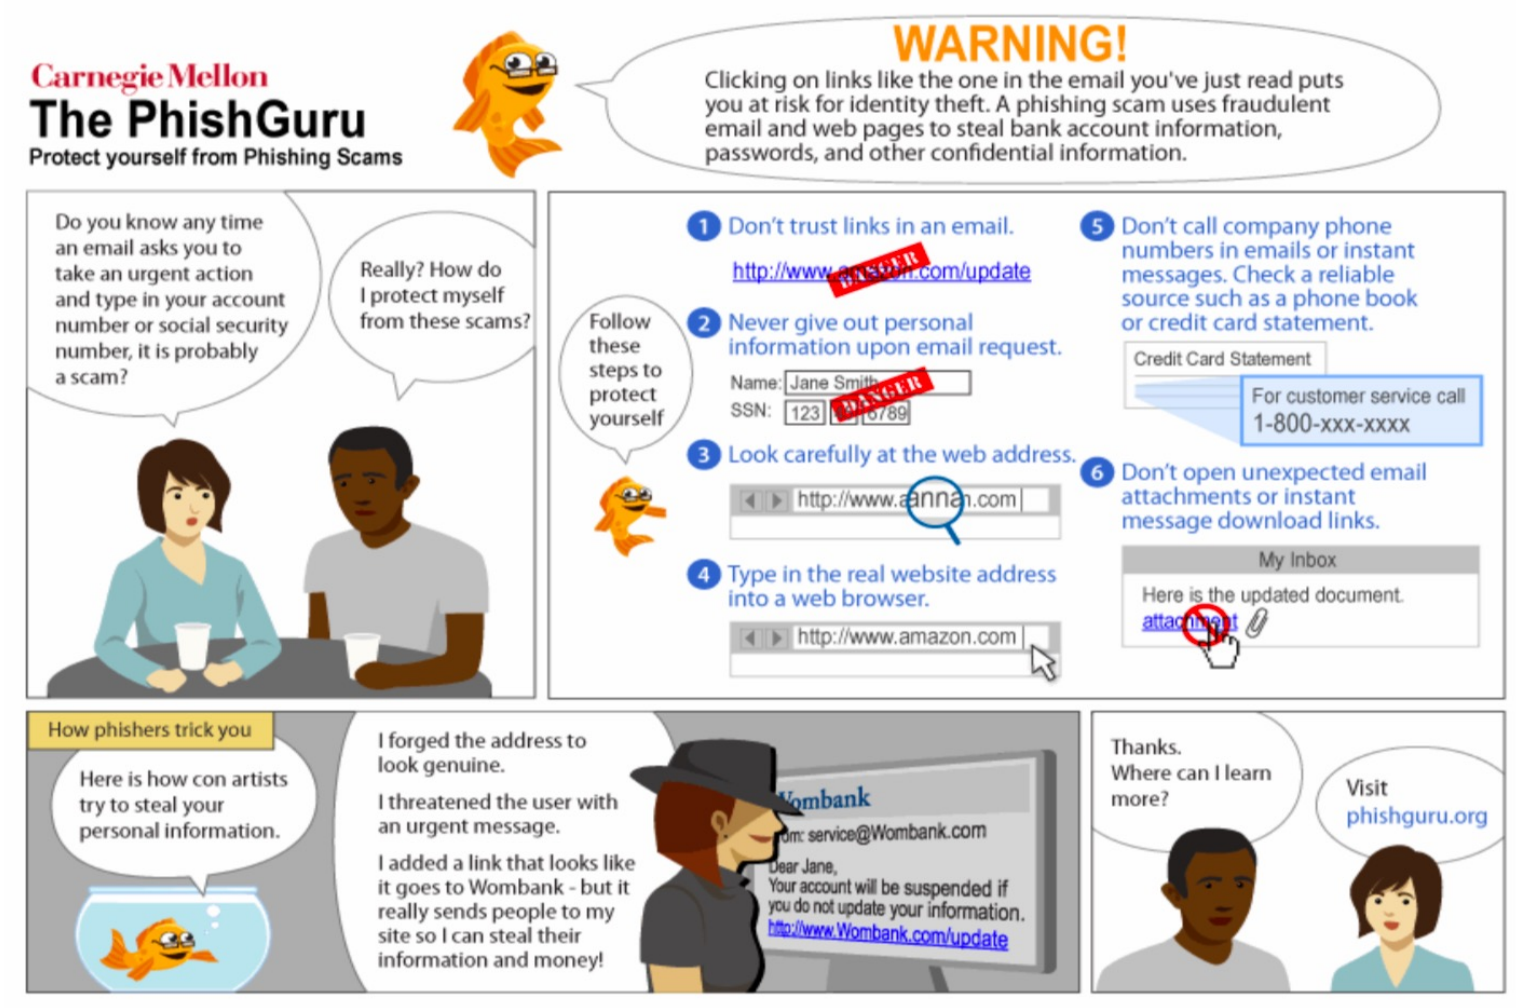

How to design and measure effective phishing training

CHI'2018

Who Provides Phishing Training? Facts, Stories, and People Like Me

Rick Wash
Michigan State University
East Lansing, MI
wash@msu.edu

Molly M. Cooper
Michigan State University and
Nova Southeastern University
East Lansing, MI
mmcooper@msu.edu

ABSTRACT
Humans represent one of the most persistent vulnerabilities in many computing systems. Since human users are independent agents who make their own choices, closing these vulnerabilities means persuading users to make different choices. Focusing on one specific human choice – clicking on a link in a phishing email – we conducted an experiment to identify better ways to train users to make more secure decisions. We compared traditional facts-and-advice training against training that uses a simple story to convey the same lessons. We found a surprising interaction effect: facts-and-advice training works better than not training users, but only when presented by a security expert. Stories don't work quite as well as facts-and-advice, but work much better when told by a peer. This suggests that the perceived origin of training materials can have a surprisingly large effect on security outcomes.

ACM Classification Keywords
K.6.5. Security and Protection: Miscellaneous

Author Keywords
phishing; training; stories; security

INTRODUCTION
interfaces with a human largely has no purpose. Humans are integral parts of modern computing systems: they provide inputs and read outputs of the system; they design, build, program, and configure the systems to work on their behalf; and they make a variety of critical decisions that the computers are unable to make themselves. In these roles, humans function as a critical component of modern computing systems, and as such, they can also create vulnerabilities in the system. For example, they may choose to click on inappropriate links in email messages they receive that install malicious software, or choose to disclose sensitive information on an inappropriate webpage (both of which go under the name "phishing").

However, unlike computers, humans cannot be programmed to perform. They cannot easily be patched to change their behavior when a vulnerability is discovered. And they do not behave deterministically, so even if they behave correctly today, they might make a different decision tomorrow. "Patching vulnerabilities" in the human users of systems is one of the most challenging aspects of computer security. Since humans are independent agents that make their own decisions, they must be *persuaded* to want to change their behavior. And most human users are not experts at using computers, and therefore often need to be *trained* to learn how to make more secure

Developed five different training materials to help teach about phishing and its harms

Facts and Advice

Phishing is an online scam involving email messages appearing to be from a trusted source. A type of phishing, called spear phishing, is especially problematic.

Spear phishing is a technique that con artists use to specifically target individuals or companies and gain access to private information or accounts.

With spear phishing, hackers disguise themselves as a trusted source by sending an email with a request to provide personal information, such as log in and password information. When the person gives the information by replying to the email or via a website link provided, the criminal goes into the account and takes what they want.

Watch for:

1. *The email urges you to take immediate action.* Often, a phishing email tries to trick you into clicking a link by claiming that your account has been closed or put on hold, or that there's been fraudulent activity requiring your immediate attention. To be safe, log into the account in question directly by visiting the appropriate website, then check your account status.
2. *The hyperlinked URL is different from the one shown.*

vs.

Personal Stories

"Sometimes frauds will target university email addresses to trick them into giving up information about themselves. I made the mistake of offering up information even after hearing this. I got a message from the "IT department" requesting that I verify my account information, otherwise my account will be suspended.

"Stupid me, I should have known that it was a trick. When I clicked on the email, it took me to a website that wasn't really my university. I had to end up canceling my account and getting a new one, changing my password, etc. It was pretty embarrassing.

"I quickly wizened up and have since never ever been a victim again. Now I hover over links to see where they link to. I won't be fooled twice."

Executed a spam campaign at their university

- Randomly targeted 2000 staff members
  - Sent **four** kinds of spam emails
  - Spread out over two months
  - Unique links for tracking which users clicked what
- If a link is clicked, user is directed to a set of (five) trainings
  - Effectiveness of a given training is determined based on if they click future phishing emails

Mailbox is almost full
New Sign-in Attempt
Upgrading Email accounts
Re-validate your mailbox

	Email No.			
Group 1	1	2	3	4
Group 2	2	4	1	3
Group 3	3	1	4	2
Group 4	4	3	2	1

Figure 2. Latin square design to counterbalance phishing emails across groups.

Ethics?

- Questions of Ecological validity
- What happens if participants knew they were part of a study?
  - Hawthorn Effect
    - Participants behave in a way to please the researchers
  - Increase likely of clicking
    - They were told they would receive benign emails

Ethical Concerns

It is difficult to conduct externally valid phishing studies ethically [14]. One of the biggest challenges is informed consent. If subjects are aware that they are participating in a research study , then they know that they may receive a fake phishing email as part of the study. This can induce two changes in the subjects. First, it can substantially reduce study validity due to changes in behavior. If people know that phishing emails they receive might be part of the study, then they might evaluate these emails differently than they normally would. This is known as "experimenter bias" or the "Hawethorne effect" – people behave differently when they know they are being watched or monitored as part of a study [28]. To avoid this bias, we did not obtain informed consent ahead of the study, following the recommendation of Finn and Jacobson [14]. Instead, we worked to minimize the potential harms of the study by ensuring that subjects did not actually encounter any negative consequences of participating in the study. For example, we did not share subject names or click rates with the university, so there were no possible employment consequences based on subject behavior.

Click rate and results…

Email	Percent Clicked
Mailbox is almost full	6.0%
New Sign-in Attempt	7.0%
Upgrading Email accounts	12.1%
Re-validate your mailbox	7.7%

Table 3. The subject lines of the four different phishing emails that we sent, and the percentage of subjects who clicked on each one. Every subject received each email exactly once.

Day	# Clicked	Percent Clicked	Repeat Clicks
0	228	11.7%	
2	143	7.35%	21.7%
7	125	6.43%	35.2%
42	146	7.51%	44.5%

Table 4. Fewer people clicked on phishing emails in later days, though there was a slight increase after a month. The last column indicates, out of the people who clicked, what percentage of them had previously clicked on the link in one of our phishing emails (and therefore previously received training).

Advice from an expert is nearly twice as effective, but…

	Clicked	Clicked Again	Percent
Advice from an Expert	76	14	18%
Control	81	19	23%
Story from a Person Like Me	89	21	24%
Advice from a Person Like Me	87	22	25%
Story from an Expert	88	30	34%

Table 7. Number of People who ever clicked after receiving training. Percentages are the same as in Table 6. $\chi^2(4, N = 421) = 6$, $p = 0.2$

$$\chi^2(4, N = 421) = 6, p = 0.2$$

No statistical difference

Spear Phishing

Spearphishing vs. Phishing

- Hyper-targetted phishing emails that are often hyper-personalized in order to gain access to an account or get money

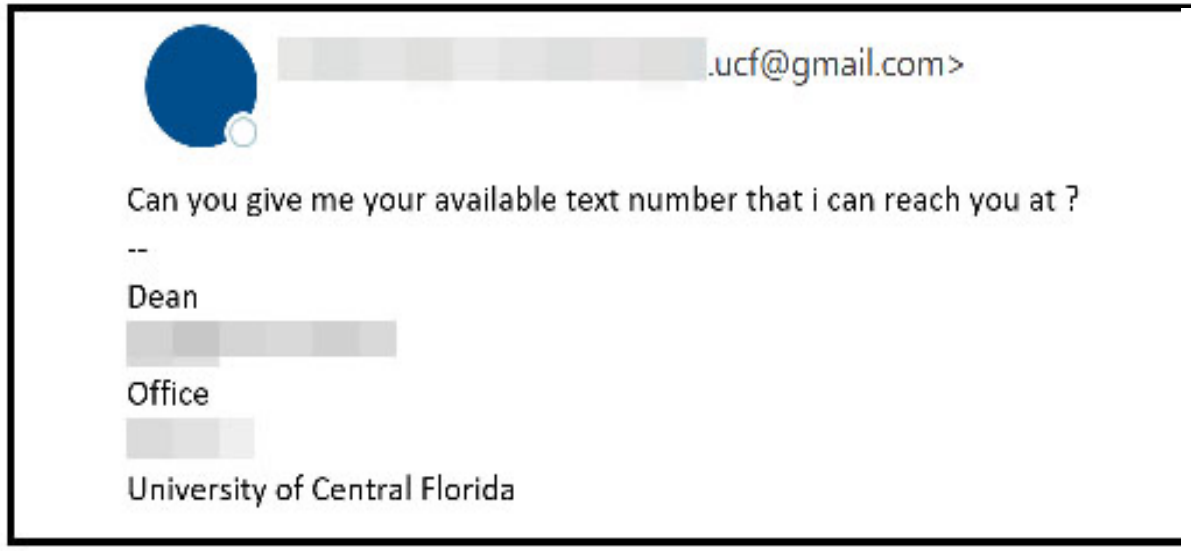

Example of a newer gift card scam asking for the recipient's perso

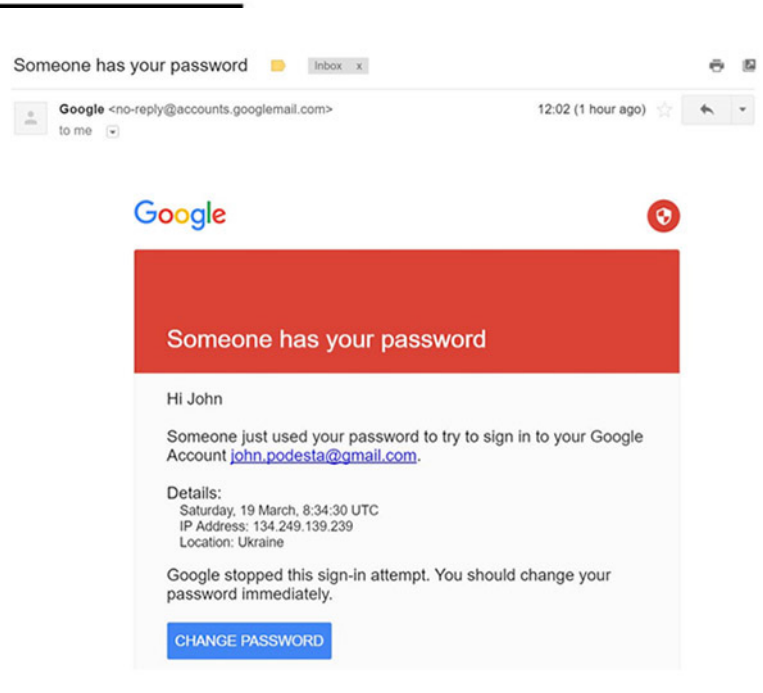

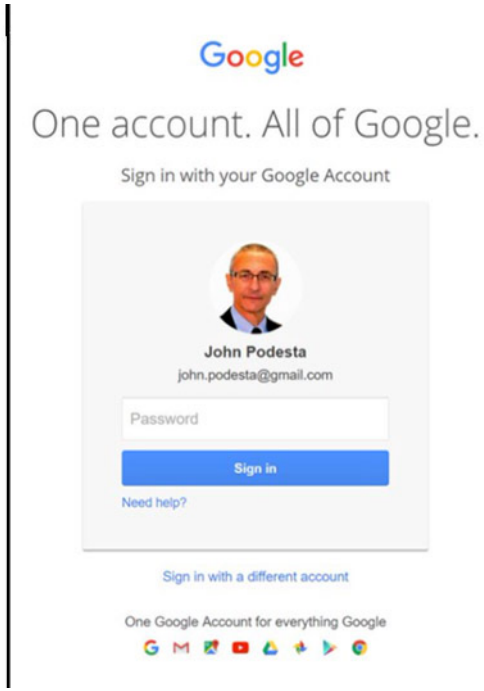

Privacy Notifications

Have you ever read a privacy policy?

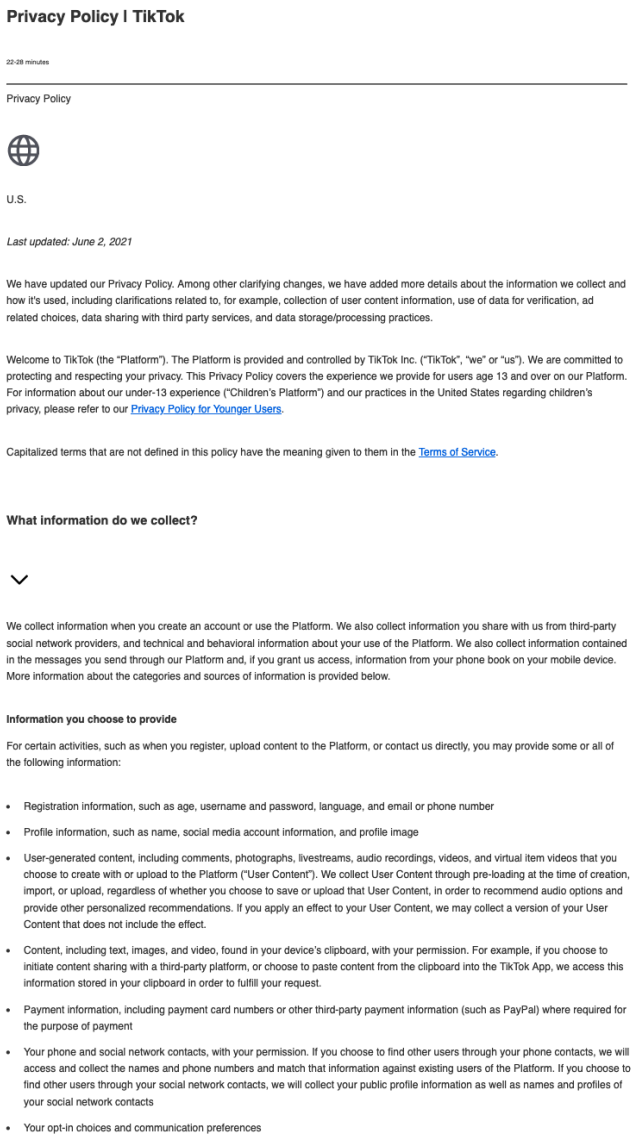

Privacy Policy | TikTok

22-28 minutes

Privacy Policy

U.S.

Last updated: June 2, 2021

We have updated our Privacy Policy. Among other clarifying changes, we have added more details about the information we collect and how it's used, including clarifications related to, for example, collection of user content information, use of data for verification, ad related choices, data sharing with third party services, and data storage/processing practices.

Welcome to TikTok (the "Platform"). The Platform is provided and controlled by TikTok Inc. ("TikTok", "we" or "us"). We are committed to protecting and respecting your privacy. This Privacy Policy covers the experience we provide for users age 13 and over on our Platform. For information about our under-13 experience ("Children's Platform") and our practices in the United States regarding children's privacy, please refer to our Privacy Policy for Younger Users.

Capitalized terms that are not defined in this policy have the meaning given to them in the Terms of Service.

What information do we collect?

We collect information when you create an account or use the Platform. We also collect information you share with us from third-party social network providers, and technical and behavioral information about your use of the Platform. We also collect information contained in the messages you send through our Platform and, if you grant us access, information from your phone book on your mobile device. More information about the categories and sources of information is provided below.

Information you choose to provide

For certain activities, such as when you register, upload content to the Platform, or contact us directly, you may provide some or all of the following information:

- Registration information, such as age, username and password, language, and email or phone number
- Profile information, such as name, social media account information, and profile image
- User-generated content, including comments, photographs, livestreams, audio recordings, videos, and virtual item videos that you choose to create with or upload to the Platform ("User Content"). We collect User Content through pre-loading at the time of creation, import, or upload, regardless of whether you choose to save or upload that User Content, in order to recommend audio options and provide other personalized recommendations. If you apply an effect to your User Content, we may collect a version of your User Content that does not include the effect.
- Content, including text, images, and video, found in your device's clipboard, with your permission. For example, if you choose to initiate content sharing with a third-party platform, or choose to paste content from the clipboard into the TikTok App, we access this information stored in your clipboard in order to fulfill your request.
- Payment information, including payment card numbers or other third-party payment information (such as PayPal) where required for the purpose of payment
- Your phone and social network contacts, with your permission. If you choose to find other users through your phone contacts, we will access and collect the names and phone numbers and match that information against existing users of the Platform. If you choose to find other users through your social network contacts, we will collect your public profile information as well as names and profiles of your social network contacts
- Your opt-in choices and communication preferences

- Not well understood
- Difficult to read
- Ineffective
- Not actively reviewed by most users

Google Privacy & Terms

Overview | Privacy Policy | Terms of Service | Technologies | FAQ

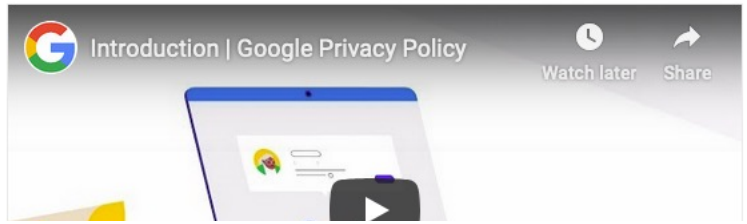

How does Google protect my privacy and keep my information secure?

We know security and privacy are important to you – and they are important to us, too. We make it a priority to provide strong security and give you confidence that your information is safe and accessible when you need it.

We're constantly working to ensure strong security, protect your privacy, and make Google even more effective and efficient for you. We spend hundreds of millions of dollars every year on security, and employ world-renowned experts in data security to keep your information safe. We also built easy-to-use privacy and security tools like Google Dashboard, 2-step verification and Ads Settings. So when it comes to the information you share with Google, you're in control.

You can learn more about safety and security online, including how to protect yourself and your family online, at the Google Safety Center.

Learn more about how we keep your personal information private and safe — and put you in control.

What kinds of information do we collect?

Data Policy

Return to top

How is this information shared?

Your informatic

Sharing on Fac

People a
When yo
audience
you selec
the public
Messenge
those pe
can also s
engagem
see who h

Public in
including
username
in your pu
Page, pub
Facebook
and we ca
off our Pr
search re
seen, acc

Sharing with Third-Party Partners
We work with third-party partners who help us provide and improve our Products or who use Facebook Business Tools to grow their businesses, which makes it possible to operate our companies and provide free services to people around the world. We don't sell any of your information to anyone, and we never will. We also impose strict restrictions on how our partners can use and disclose the data we provide. Here are the types of third parties we share information with:

Partners who use our analytics services.
We provide aggregated statistics and insights that help people and businesses understand how people are engaging with their posts, listings, Pages, videos and other content on and off the Facebook Products. For example, Page admins and Instagram business profiles receive information about the number of people or accounts who viewed, reacted to, or commented on their posts, as well as aggregate demographic and other information that helps them understand interactions with their Page or account.

Advertisers.
We provide advertisers with reports about the kinds of people seeing their ads and how their ads are performing, but we don't share information that personally identifies you (information such as your name or email address that by itself can be used to contact you or identifies who you are) unless you give us permission. For example, we provide general demographic and interest information to advertisers (for example, that an ad was seen by a woman between the ages of 25 and 34 who lives in Madrid and likes software engineering) to help them better understand their audience. We also confirm which Facebook ads led you to make a purchase or take an action with an advertiser.

Content others share or reshare about you
You should consider who you choose to share with, because people who can see your activity on our Products can choose to share it with others on and off our Products, including people and businesses outside the

Measurement partners.
We share information about you with companies that aggregate it to provide analytics and measurement reports to our partners.

Partners offering goods and services in our Products.
When you subscribe to receive premium content, or buy something from a seller in our Products, the content creator or seller can receive your public information and other information you share with them, as well as the information needed to complete the transaction, including shipping and contact details.

Vendors and service providers.
We provide information and content to vendors and service providers who support our business, such as by providing technical infrastructure services, analyzing how our Products are used, providing customer service, facilitating payments or conducting surveys.

Researchers and academics.
We also provide information and content to research partners and academics to conduct research that advances scholarship and innovation that support our business or mission, and enhances discovery and innovation on topics of general social welfare, technological advancement, public interest, health and well-being.

Law enforcement or legal requests.
We share information with law enforcement or in response to legal requests in the circumstances outlined below.

Apps, websites, and third-party integrations on or using our Products.
When you choose to use third-party apps, websites, or other services that use, or are integrated with, our Products, they can receive
u post or share. For example, when you play a
friends or use a Facebook Comment or Share
ame developer or website can receive
tivities in the game or receive a comment or
e website on Facebook. Also, when you
d-party services, they can access your public
ny information that you share with them. Apps
receive your list of Facebook friends if you
m. But apps and websites you use will not be
nformation about your Facebook friends from
any of your Instagram followers (although your
of course, choose to share this information
collected by these third-party services is subject
licies, not this one.

tems providing native versions of Facebook and
ave not developed our own first-party apps)
mation you choose to share with them,
friends share with you, so they can provide
ou.

ss of restricting developers' data access even
use. For example, we will remove developers'
nd Instagram data if you haven't used their app
anging Login, so that in the next version, we
n app can request without app review to
am username and bio, profile photo and email
ther data will require our approval.

Problems with Self Management

- Uninformed Individual
  - “notice and choice” – but most people don’t read privacy policies
  - **Privacy notices are long and difficult**
  - Users operated under incorrect assumptions and are not well informed about how their data could be used
  - *Unmotivated user?*
- Skewed Decision Making
  - Users cannot assess the consequences of disclosure
  - Disconnect between high value of privacy and behavior
    - **Privacy paradox!**

Structural Problems with Self Management

- Scale
  - A user may be able to manage one or two entities, but not at the scale of the internet
- Aggregation
  - Sharing in isolation may be combined to create aggregate profiles
  - Containing privacy in one entity may not affect others
- Assessing Harm
  - Immediate benefits outweigh “future detriments”
  - Consent is a short-term choice, disclosure is a long-term effect

Understanding the space of notices (c. 2015)

- Survey existing privacy notices
- Identified challenges, requirements, and best practices in the design of privacy notices
- Systemization of knowledge and a taxonomy

SOUPS 2015

A Design Space for Effective Privacy Notices

Florian Schaub,[1] Rebecca Balebako,[2]* Adam L. Durity,[3]* Lorrie Faith Cranor[1]

[1]Carnegie Mellon University
Pittsburgh, PA, USA
{ fschaub, lorrie }@cmu.edu

[2]RAND Corporation
Pittsburgh, PA, USA
balebako@rand.org

[3]Google
Mountain View, CA, USA
adurity@google.com

ABSTRACT

Notifying users about a system's data practices is supposed to enable users to make informed privacy decisions. Yet, current notice and choice mechanisms, such as privacy policies, are often ineffective because they are neither usable nor useful, and are therefore ignored by users. Constrained interfaces on mobile devices, wearables, and smart home devices connected in an Internet of Things exacerbate the issue. Much research has studied usability issues of privacy notices and many proposals for more usable privacy notices exist. Yet, there is little guidance for designers and developers on the design aspects that can impact the effectiveness of privacy notices. In this paper, we make multiple contributions to remedy this issue. We survey the existing literature on privacy notices and identify challenges, requirements, and best practices for privacy notice design. Further, we map out the design space for privacy notices by identifying relevant dimensions. This provides a taxonomy and consistent terminology of notice approaches to foster understanding and reasoning about notice options available in the context of specific systems. Our systemization of knowledge and the

website, or linked to from mobile app stores or mobile apps, to signs posted in public places to inform about CCTV cameras in operation. Even an LED indicating that a camera or microphone is active and recording constitutes a privacy notice, albeit one with limited information about the data practices associated with the recording. Providing notice about data practices is an essential aspect of data protection frameworks and regulation around the world [57]. While transparency has been emphasized as an important practice for decades, existing privacy notices often fail to help users make informed choices. They can be lengthy or overly complex, discouraging users from reading them.

Smartphones and mobile apps introduce additional privacy issues as they support recording of sensor and behavioral information that enables inference of behavior patterns and profiling of users. Yet, comparatively smaller screens and other device restrictions constrain how users can be given notice about and control over data practices.

The increasing adoption of wearable devices, such as smart watches or fitness trackers, as well as smart home devices, such as smart thermostats, connected light bulbs, or smart

Challenges for Effective Privacy Notices

- Notice Complexity
  - Conflation of requirements
  - Long privacy policies full of jargon, legal vocab, business relations
  - Purposefully vague and unclear to regular users
- Lack of Choices
  - Inform but offer no real choices
  - Using the service, regardless of reading the notice is interpreted as consent
  - Can be mere warnings: “CCTV in use”
- Notice Fatigue
  - Feel like it is pointless to read them because of complexity and a lack of choice
  - Businesses may change privacy notices frequently, leading to more notice and dismissal at a high frequency
- Decoupled Noticed
  - Notice is presented in one context, e.g., on a website or a manual, but not in the larger context in which the data may be used
  - May not read the notice for the right context or service

How much time and money does it cost to stay up to date on privacy policies? (c. 2008)

- Online study with 212 participants to measure the time to skim online privacy policies
- Reading time estimates based on 75 most popular websites, average reading rate of 250 words per minute
- Complete Reading time
  - Lower bound 181 hours
  - Upper bound 304 hours
  - Estimate: 304 hours
- Skimming Reading Time
  - Lower Bound 81 hours
  - Upper Bound 154 hours
  - Estimate: 293 hours

- *"We present a range of values, and found the national opportunity cost for just the time to read policies is on the order **of $781** billion"*.

I/S: Journal of Law and Policy **2008**

The Cost of Reading Privacy Policies

ALEECIA M. MCDONALD & LORRIE FAITH CRANOR*

Abstract: Companies collect personally identifiable information that website visitors are not always comfortable sharing. One proposed remedy is to use economics rather than legislation to address privacy risks by creating a marketplace for privacy where website visitors would choose to accept or reject offers for small payments in exchange for loss of privacy. The notion of micropayments for privacy has not been realized in practice, perhaps because advertisers might be willing to pay a penny per name and IP address, yet few people would sell their contact information for only a penny.[1] In this paper we contend that the time to read privacy policies is, in and of itself, a form of payment. Instead of receiving payments to reveal information, website visitors must pay with their time to research policies in order to retain their privacy. We pose the question: if website users were to read the privacy policy for each site they visit just once a year, what would their time be worth?

REPORT TO THE PRESIDENT

BIG DATA AND PRIVACY: A TECHNOLOGICAL PERSPECTIVE

4.3 Notice and consent

Notice and consent is, today, the most widely used strategy for protecting consumer privacy. When the user downloads a new app to his or her mobile device, or when he or she creates an account for a web service, a notice is displayed, to which the user must positively indicate consent before using the app or service. In some fantasy world, users actually read these notices, understand their legal implications (consulting their attorneys if necessary), negotiate with other providers of similar services to get better privacy treatment, and only then click to indicate their consent. Reality is different.[118]

Privacy Notice Design Space

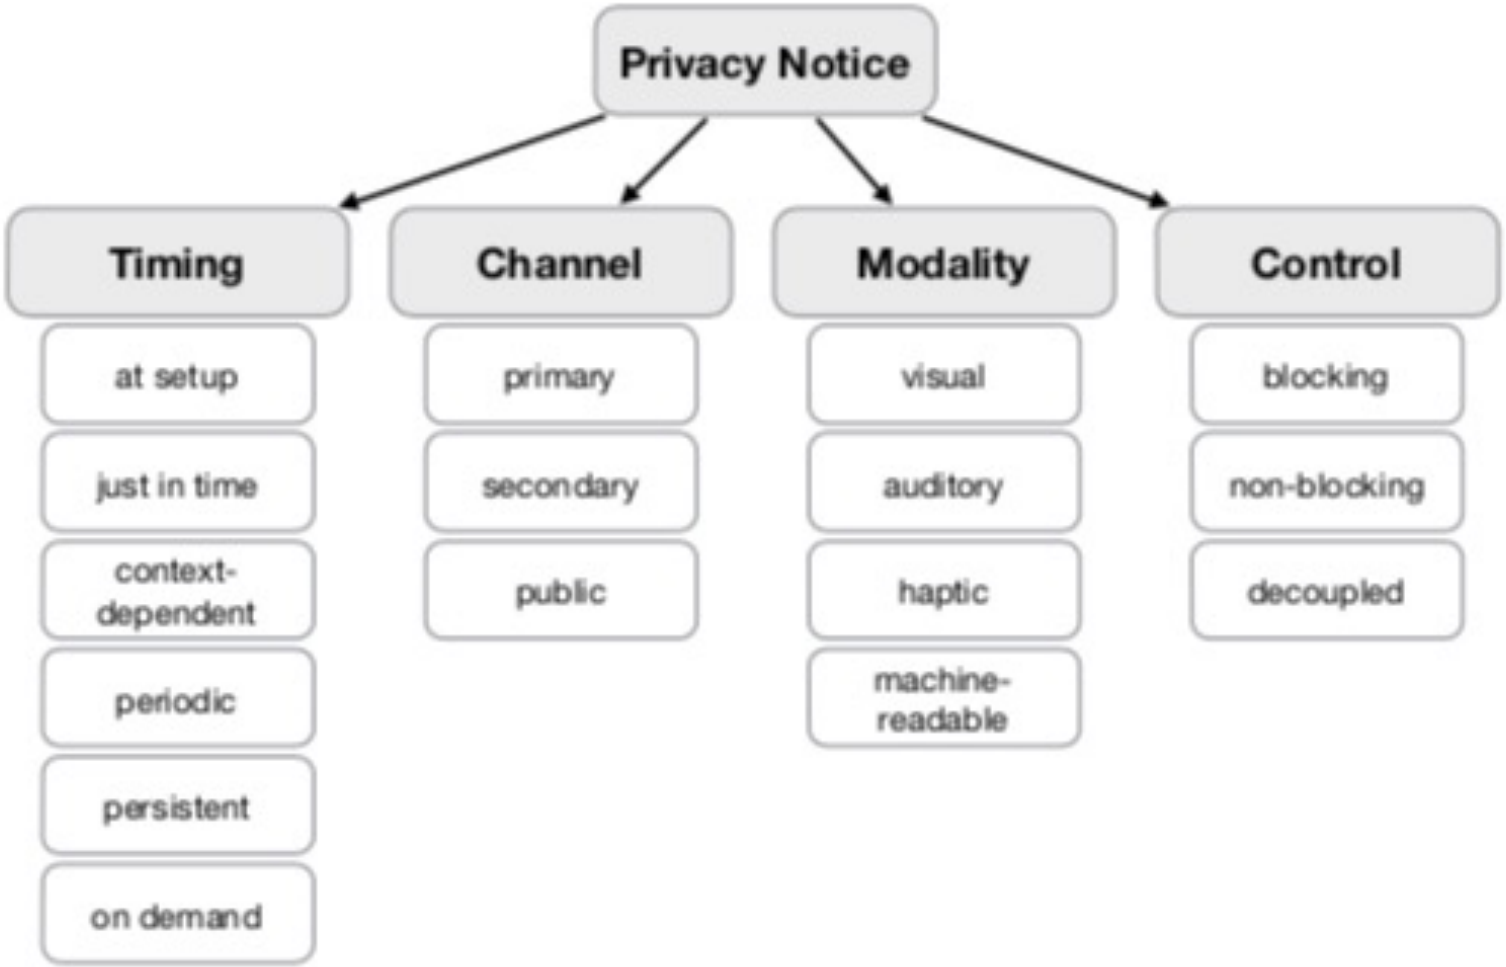

Timing of Notices

- At setup
  - Used fir the first time
- Just in time
  - When active vs. inactive
- Context dependent
  - When changes occur, such as a change in location
- Periodic
  - How frequently after initial setup do notices keep occurring?
- Persistent
  - Show notice all the time
- On Demand
  - “privacy dashboard”

Channel – delivery mechanism of notice

- Primary
  - Provided on the same platform of the device
- Secondary
  - Provided on second channel
  - E.g., wearables, devices without displays
- Public
  - Public notices when specific users may be hard to identify

How are privacy notice displayed and communicated?

- Visual
  - Text notices
  - Images and icons
- Auditory
  - Spoken word
  - Sounds
    - Error/Warning message
- Haptic
  - E.g., Vibration on wearable
- Machine Readable
  - DoNotTrack, Cookies
  - Within the protocotols

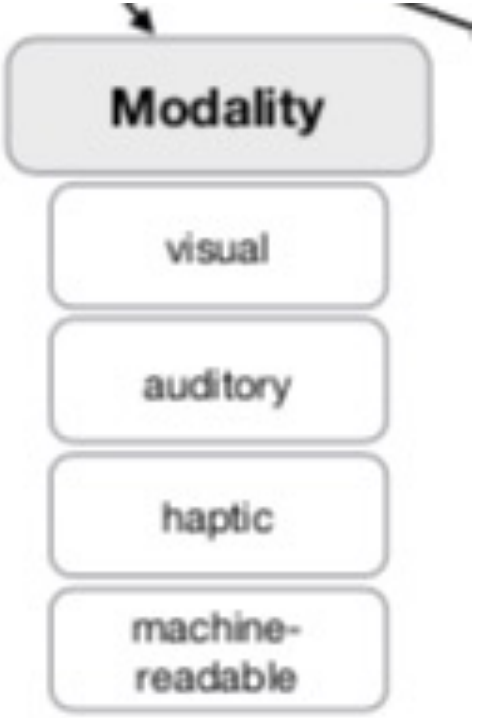

How does the service respond to consent?

- Blocking
  - Cannot continue until providing consent
  - Cannot use service without consent
- Non-Blocking
  - Can limit access and still use the service
  - E.g., sharing settings on FB
- Decoupled
  - No direct choice of access, but informs users of places to observe and alter privacy later
  - E.g., privacy dashboards

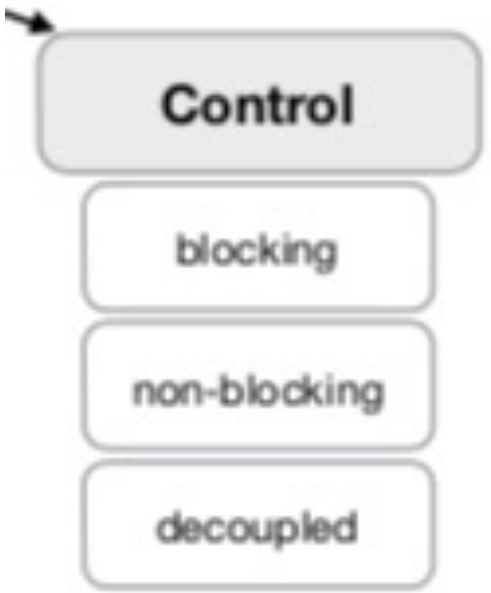

Example

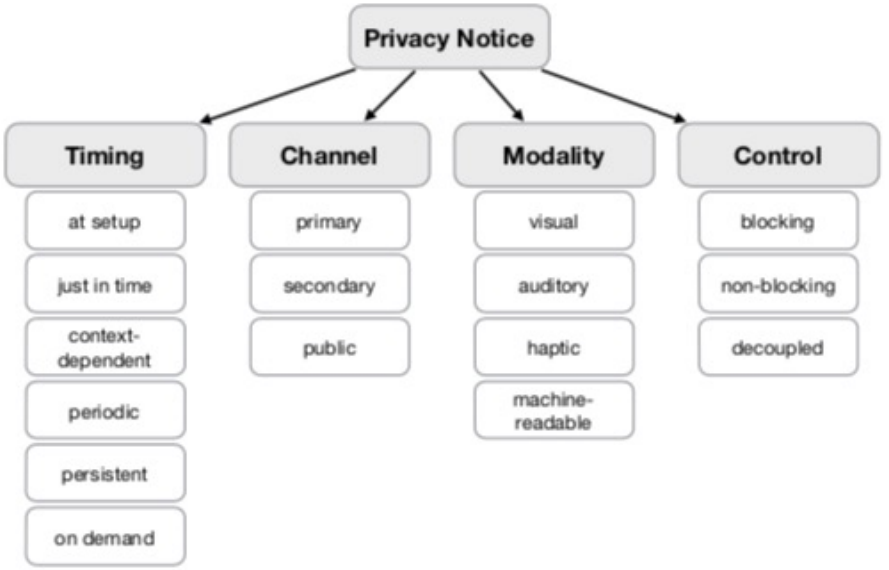
Privacy Notice
Timing
at setup
just in time
context-dependent
periodic
persistent
on demand
Channel
primary
secondary
public
Modality
visual
auditory
haptic
machine-readable
Control
blocking
non-blocking
decoupled


Privacy Checkup
Skip
Hi Charlie — Sorry to interrupt. You haven't changed who can see your posts lately, so we just wanted to make sure you're sharing this post with the right audience. (Your current setting is Public, though you can change this whenever you post.) Learn more.
Who do you want to share this post with?
Friends
Public
More Options

Example

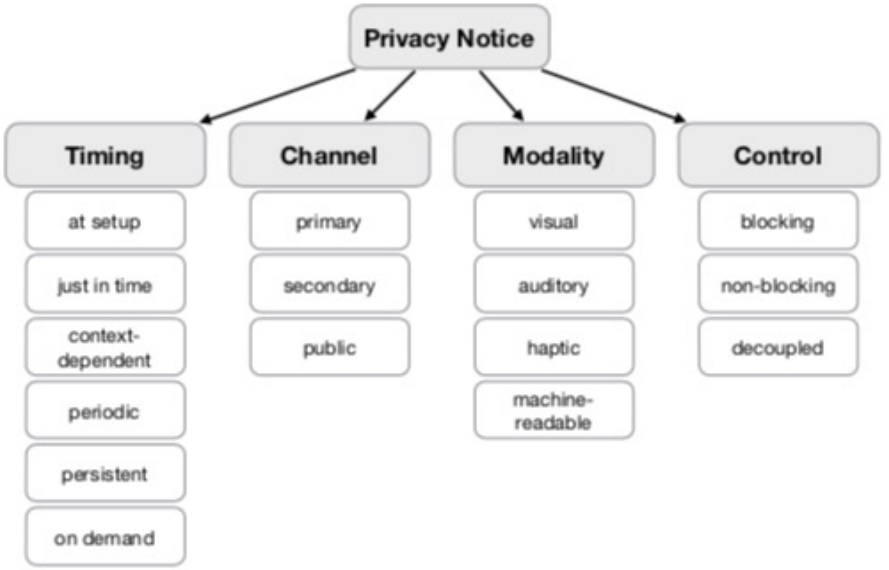
Privacy Notice
Timing
Channel
Modality
Control
at setup
just in time
context-dependent
periodic
persistent
on demand
primary
secondary
public
visual
auditory
haptic
machine-readable
blocking
non-blocking
decoupled

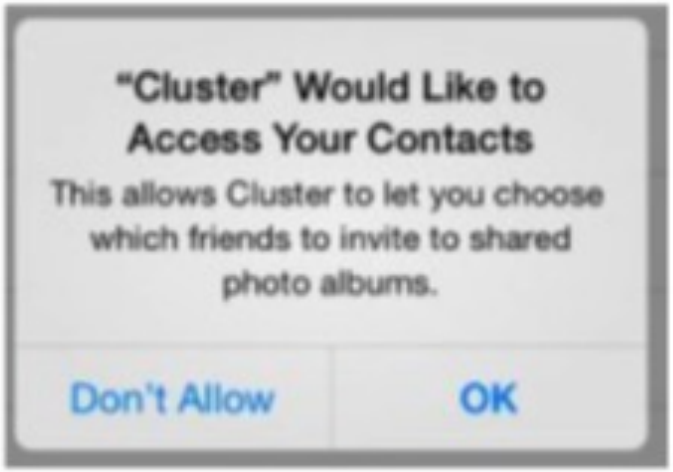
"Cluster" Would Like to Access Your Contacts
This allows Cluster to let you choose which friends to invite to shared photo albums.
Don't Allow
OK

Example

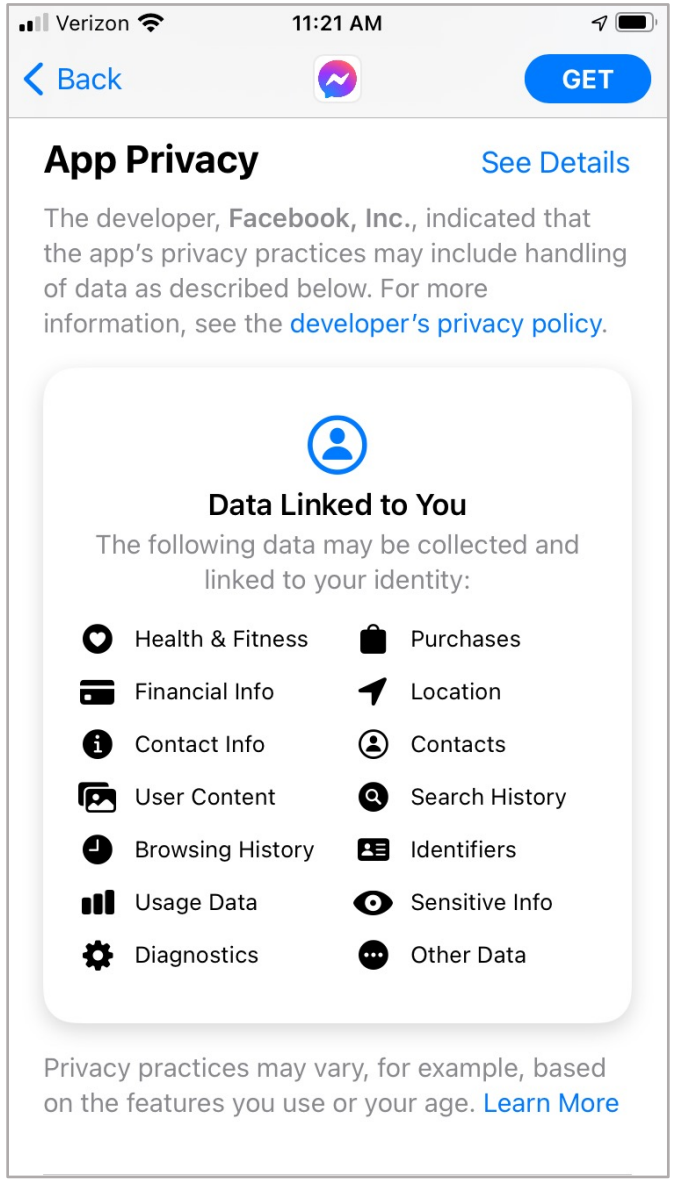
Verizon
11:21 AM
Back
GET
App Privacy
See Details
The developer, Facebook, Inc., indicated that the app's privacy practices may include handling of data as described below. For more information, see the developer's privacy policy.
Data Linked to You
The following data may be collected and linked to your identity:
Health & Fitness
Purchases
Financial Info
Location
Contact Info
Contacts
User Content
Search History
Browsing History
Identifiers
Usage Data
Sensitive Info
Diagnostics
Other Data
Privacy practices may vary, for example, based on the features you use or your age. Learn More

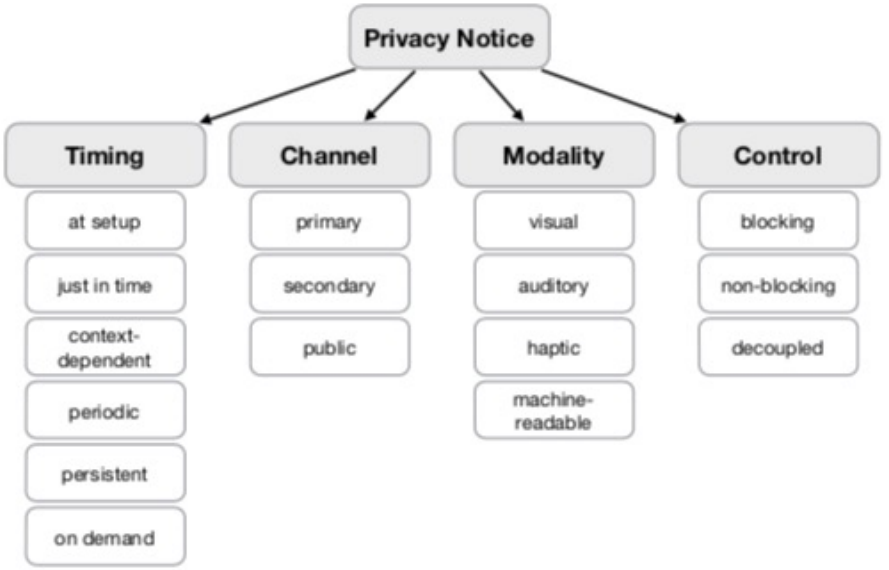
Privacy Notice
Timing
at setup
just in time
context-dependent
periodic
persistent
on demand
Channel
primary
secondary
public
Modality
visual
auditory
haptic
machine-readable
Control
blocking
non-blocking
decoupled

Other examples?

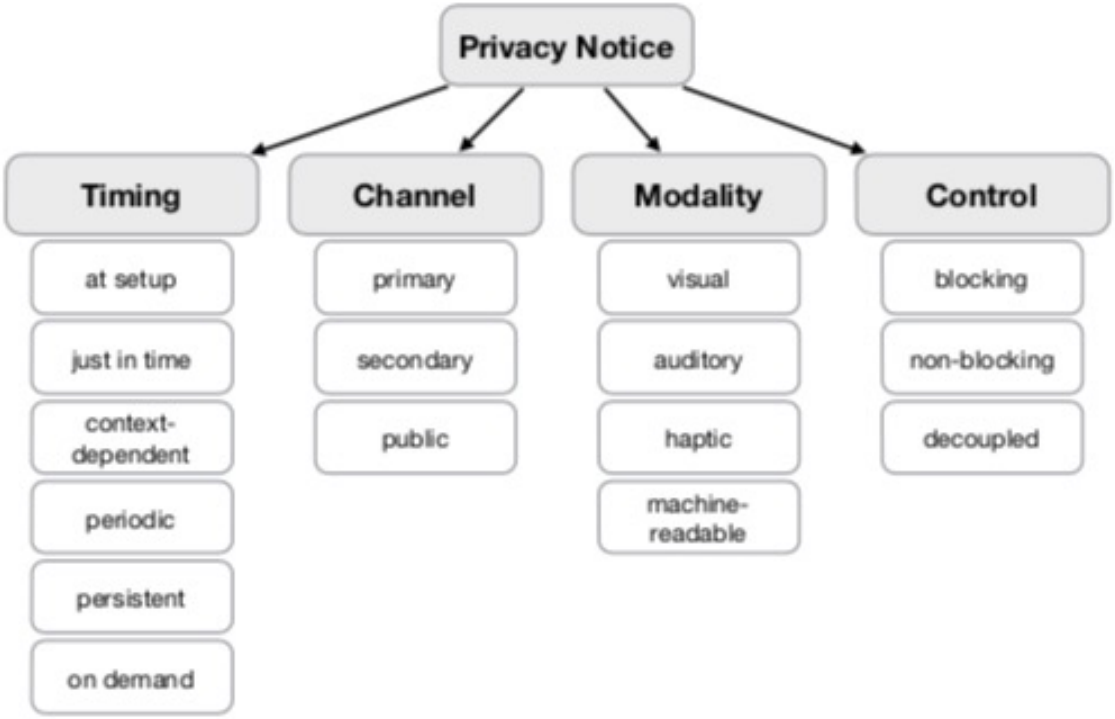

Discussion Questions

- How are we doing on privacy notices today?
- Are there features of the design space that are missing?
- Can we use the design space to evaluate other kinds of notices, or is it specific to privacy?
- Have you ever used a privacy dashboard?

`https://myactivity.google.com`