Combating Phishing Attacks: A Never Ending Arms Race?

Norman Sadeh
School of Computer Science
Carnegie Mellon University

Work conducted in collaboration with:
Lorrie Cranor, Jason Hong, Alessandro Acquisti, Julie Downs,
Sven Dietrich, Anthony Tomasic, and many graduate students

Carnegie Mellon CyLab

Carnegie Mellon School of Computer Science

Outline

- What is Phishing?
- How big of a problem is it?
- What are the current solutions?
- Carnegie Mellon University's research on "Supporting Trust Decisions"
  - Largest US research projects on combating phishing attacks (2005-present)
  - Multi-pronged approach
  - Studies and existing deployments
- Concluding Remarks
- Q&A

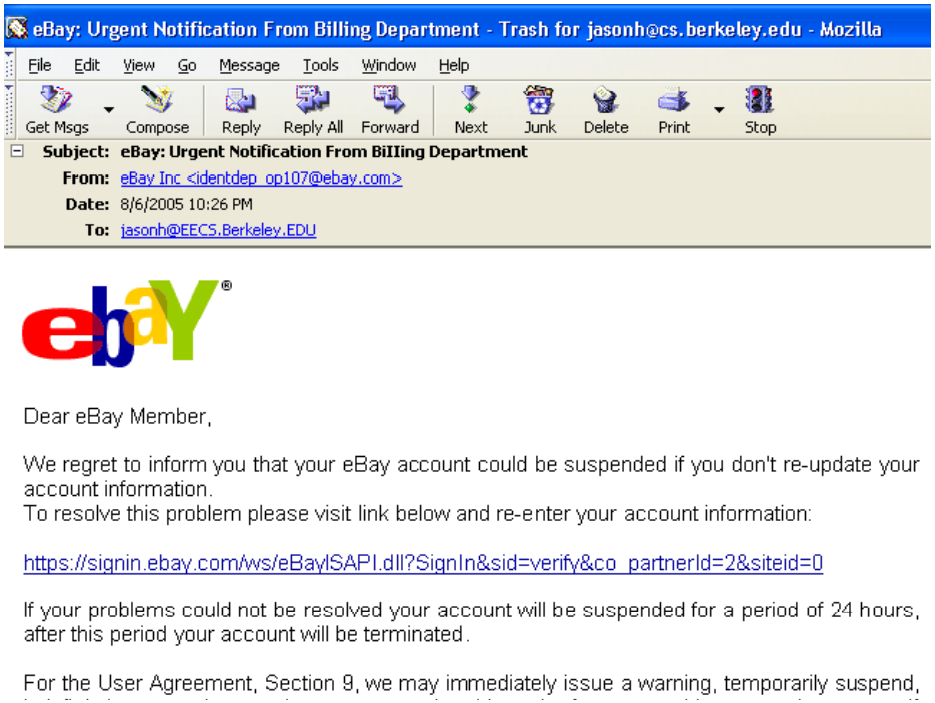

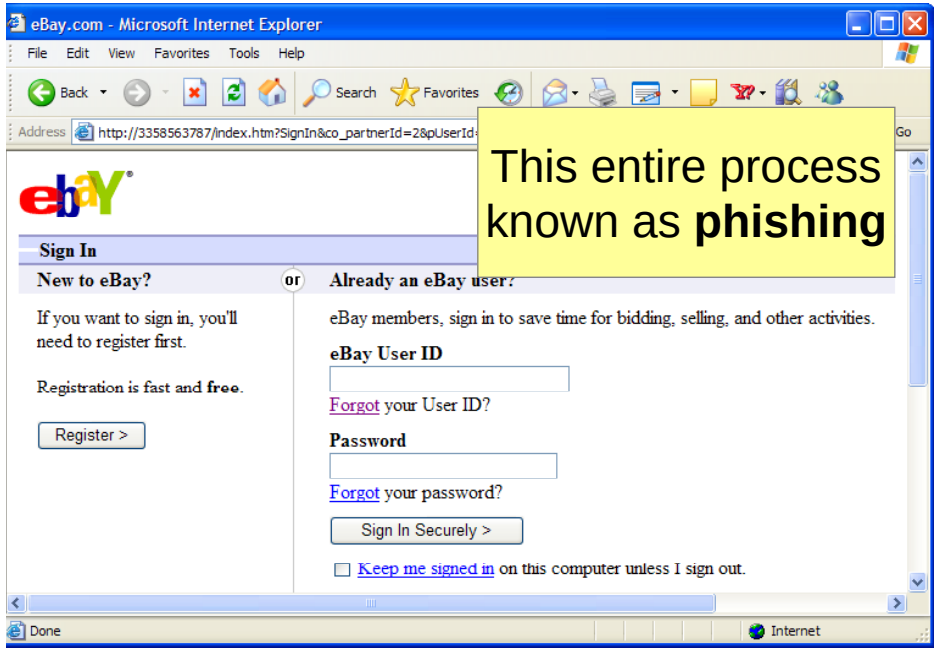

This entire process known as **phishing**

Phishing is a Plague on the Internet

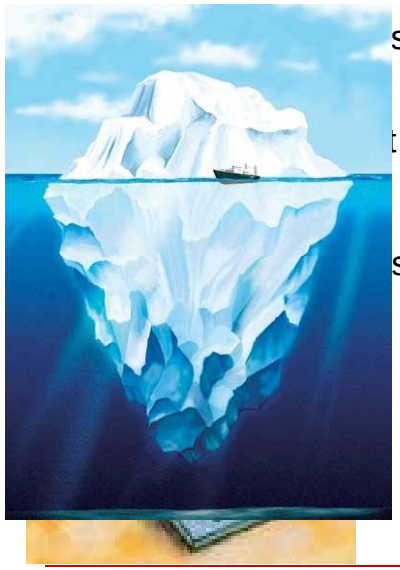

shing sites reported in Nov

t losses a year (sources: Google &

damage to brand, sales, etc.

...Just the tip of the iceberg...

...Spear Phishing Attacks...

Spear Phishing

- Phishing attacks also used for espionage
  - Government (e.g. IRS, Medicare/Medicaid, DoD)
  - Corporate, incl. pharmaceutical
    - Recent salesforce.com phish
  - Recent trend in targeting executives

...and More...

- SMiShing – SMS phishing
  - Example: "We're confirming you've signed up for our da
    unless
    - Wh
      inst
- Vishing -
  - Exploi
  - Instru
    perso
- Social ne

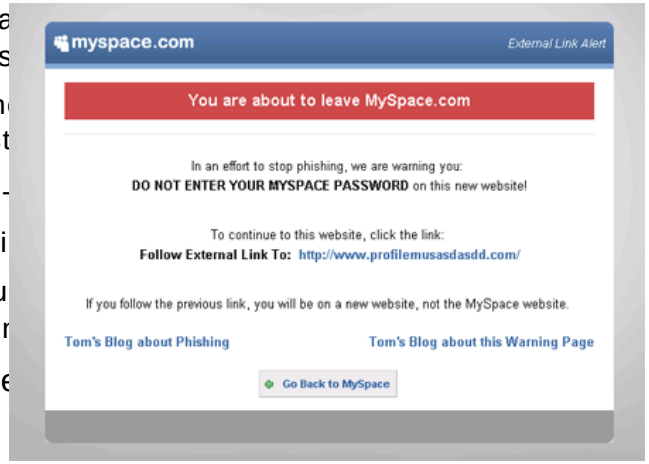

Why do people fall for phishing attacks?

Interview Study

- Interviewed 40 Internet users, including 35 non-experts
- "Mental models" interviews included email role play and open ended questions
- Interviews recorded and coded

J. Downs, M. Holbrook, and L. Cranor.
Decision Strategies and Susceptibility to Phishing.
In *Proc. of the 2006 Symposium On Usable Privacy and Security*

Little knowledge of phishing

- Only about half knew meaning of the term "phishing"

"Something to do with the band Phish, I take it."

Minimal knowledge of lock icon

- 85% of participants were **aware of lock** icon
- Only **40%** of those knew that it was **supposed to be in the browser chrome**
- **Only 35% had noticed https**, and many of those did not know what it means

"I think that it means secured."

"It symbolizes some kind of security, somehow."

Little attention paid to URLs

- **Only 55%** of participants said they had ever **noticed an unexpected or strange-looking URL**
- Most did **not** consider them to be **suspicious**

"If it wasn't one of my standard dot-com, dot-edu, dot-us or some country code, then I would be really curious what that meant."

Some knowledge of scams

- **55%** of participants reported being **cautious** when email asks for **sensitive financial** info
- But very **few reported being suspicious** of email asking for **passwords**
- **Knowledge** of financial phish **reduced likelihood** of falling for these scams
- But those knowledgeable about financial phish were **not necessarily suspicious of other scams**, such as amazon.com password phish

Naive evaluation strategies

- The most frequent strategies don't help much in identifying phish
  - This email appears to be for me
  - It's normal to hear from companies you do business with
  - **Reputable companies will send emails**

"I will probably give them the information that they asked for. And I would assume that I had already given them that information at some point so I will feel comfortable giving it to them again."

Past experience helps some

- Those who had seen scams in the past could identify similar scams
  - All but one participant correctly identified a Katrina email message as a scam
- But **knowledge of some scams didn't seem to help them identify other types of scams**

Today's Solutions Fall Short

Today's Solutions - in a Nutshell

- **Public Key Infrastructure or multi-factor authentication**
  - People don't understand certificates
  - Secure sign-on, using pictures and passphrases
- **Usability:** people don't notice/ignore

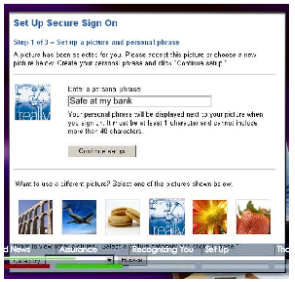

Copyright © Lorrie Cranor, Jason Hong & Norman Sadeh, 2006-2008

Today's Solutions - in a Nutshell

- **Anti-phishing filters that rely on blacklists and whitelists**
  - Always one step behind & possible liability

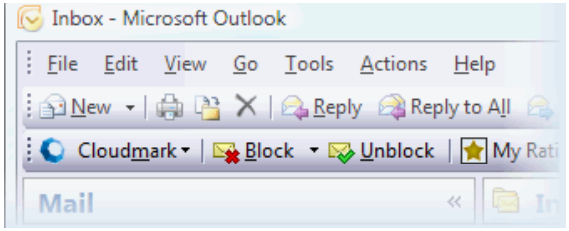

Copyright © Lorrie Cranor, Jason Hong & Norman Sadeh, 2006-2008

Today's Solutions - in a Nutshell

- **Some training** – e.g. websites, posters
  - Users just don't pay attention

Copyright © Lo

Today's Solutions - in a Nutshell

- **Spam filters** have limited success catching phish
  - In contrast to spam, phishing emails look legitimate

What else can we do?

Multi-Pronged Approach

- **Human side**
  - Interviews to understand decision-making
  - Embedded training
  - Anti-phishing game
- **Computer side**
  - Email anti-phishing filter
  - Automated testbed for anti-phishing toolbars
  - Anti-phishing toolbar
- **Automate where possible, support where necessary**

PILFER Email Filter

- **Rationale**: Spam filters let a large number of phishing emails slip through
- **Solution**: **Phishing email filter** combining a set of features aimed at catching deception along with **advanced machine learning** techniques
- Can work in **standalone** mode **or as complement to spam filter**
  - e.g. available as SpamAssassin plugin

I. Fette, N. Sadeh, and A. Tomasic
"Learning to Detect Phishing Emails",
Proceedings of the 16th International
World Wide Web Conference, May 2007 (WWW2007):

PILFER: Approach

- **Features intended to detect deception**
  - Emails are made to look like legitimate emails from the company, so text analysis is of limited value
- **Included features**:
  - Age of linked-to domains
  - Number of domains linked to
  - Presence of attention-directing links ("click here") that link to a domain other than the most common one in the email
  - ..
- Trained using **Random Forests**

PILFER Performance

- Performance on corpus of 90,000 emails

	False Positives	False Negatives
PILFER	0.13%	4.79%
SpamAssassin	3.19%	7.68%

PILFER has **25x fe**

PILFER is **20 times faster** too!

CANTINA Web Page Filter

- **Rationale**: Blacklists fall short
  - Always one step behind
  - Only protect from most common attacks, i.e. no spearphishing protection
  - Easily overcome by criminals
- **Solution**: Uses a content-based approach
  - Lexical signature & PageRank to identify phishing sites
  - Highly effective
  - No human intervention required
  - Protects against spearphishing

Y. Zhang, J. Hong, and L. Cranor.
CANTINA: A content-based approach to detecting phishing web sites.
In *Proc. of the 16th Intl. conf. on World Wide Web*, 2007 (WWW2007).

CANTINA: Phishing Web Site Detector

- CANTINA uses a simple **content-based approach**
  - Examines content of a web page and creates a "**fingerprint**"
    - Uses TF/IDF
  - Sends that fingerprint as a query to a **search engine**
  - Sees if the web page in question is in the top search results
    - If so, then we label it legitimate
    - Otherwise, we label it phishing
  - Some additional heuristics
- **Nice properties**:
  - Fast
  - Scales well
  - No maintenance by us (done by search engines)
  - Highly accurate

Evaluating CANTINA (Iteration #2)

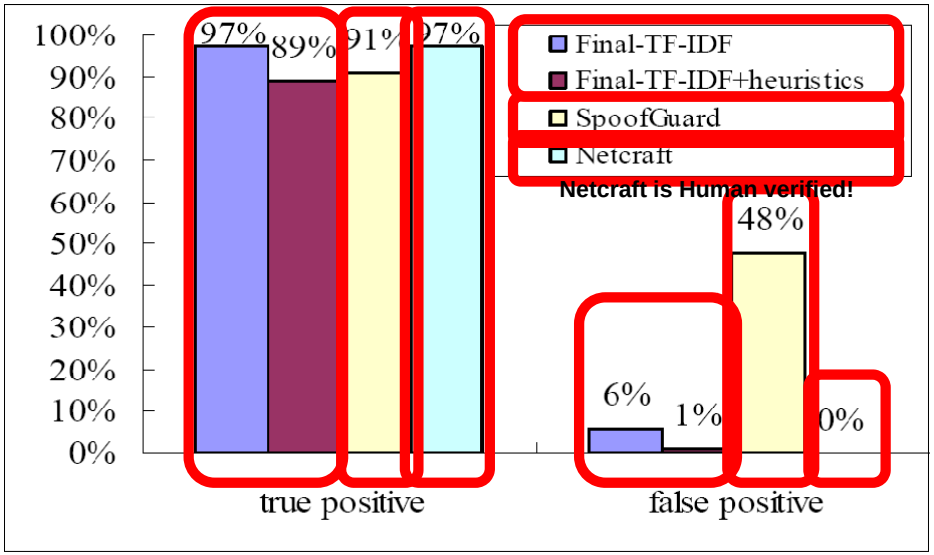

Anti-Phishing Phil Game

- **Rationale**: Traditional training doesn't work
  - But people do like playing games
- **Solution**: A game teaching about phishing

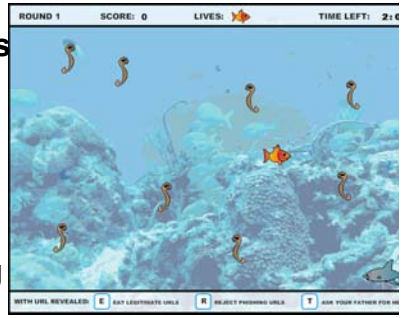

Results with **tens of thousands of users** show:

- People more **willing to play** game than read training
- People **better at identifying phishing sites after playing**

S. Sheng, B. Magnien, P. Kumaraguru, A. Acquisti, L. Cranor, J. Hong, and E. Nunge. Anti-Phishing Phil: The Design and Evaluation of a Game That Teaches People Not to Fall for Phish. In *Proc. of the 2007 Symposium On Usable Privacy and Security*, 2007.

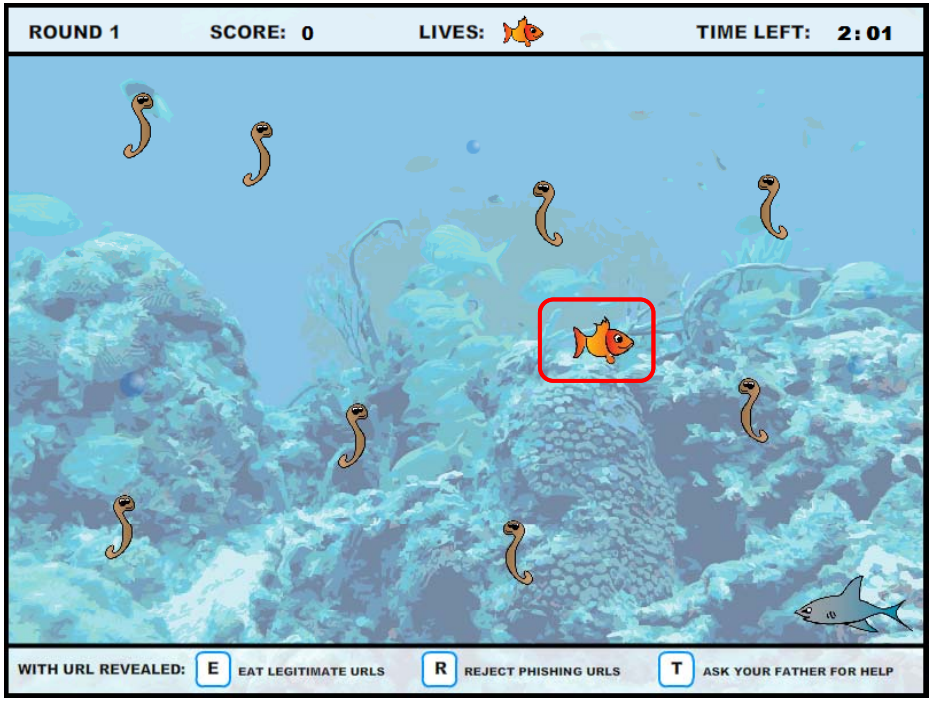

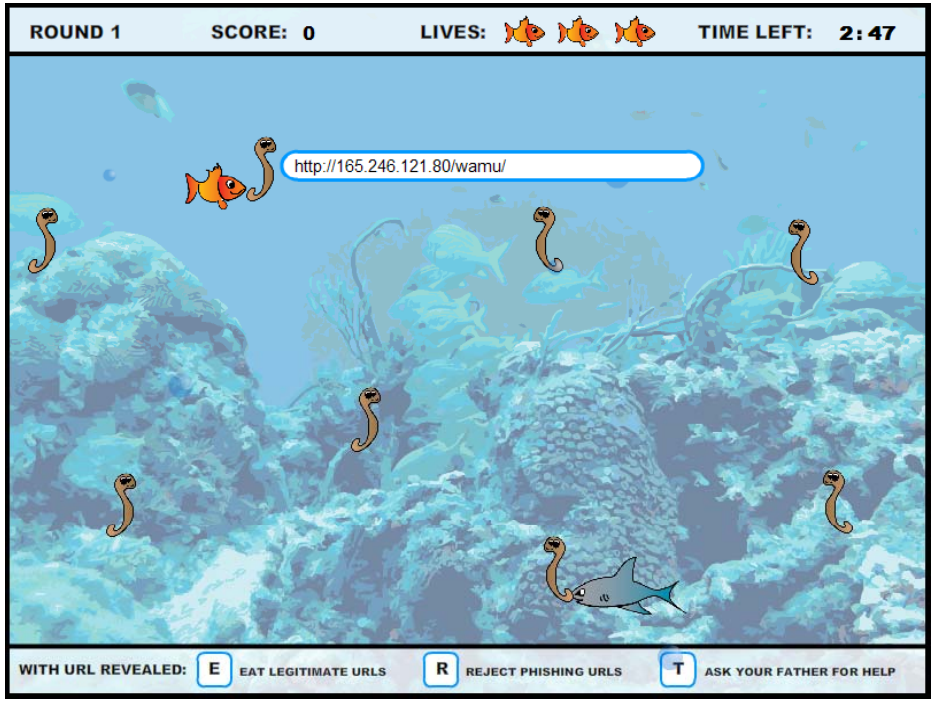

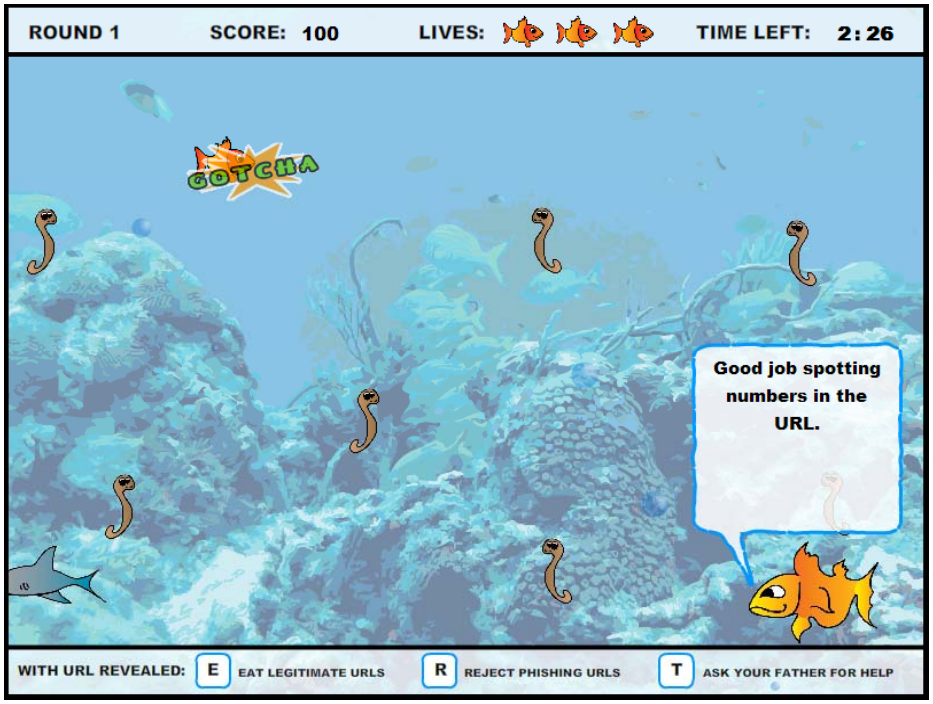
ROUND 1
SCORE: 100
LIVES:
TIME LEFT: 2:26
GOTCHA
Good job spotting numbers in the URL.
WITH URL REVEALED:
E EAT LEGITIMATE URLS
R REJECT PHISHING URLS
T ASK YOUR FATHER FOR HELP

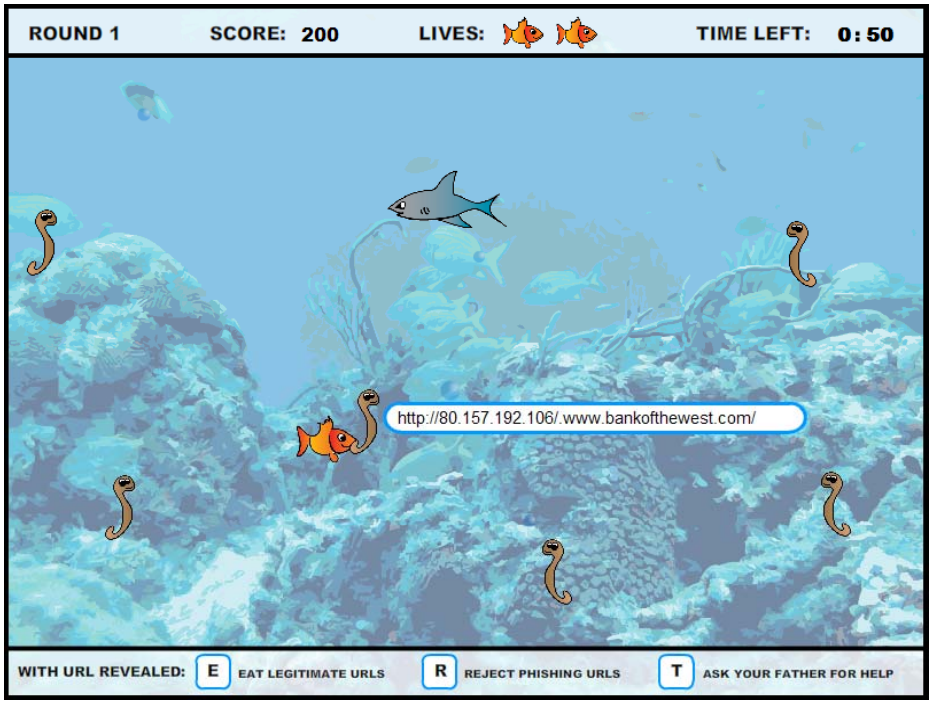
ROUND 1
SCORE: 200
LIVES:
TIME LEFT: 0:50
http://80.157.192.106/.www.bankofthewest.com/
WITH URL REVEALED:
E EAT LEGITIMATE URLS
R REJECT PHISHING URLS
T ASK YOUR FATHER FOR HELP

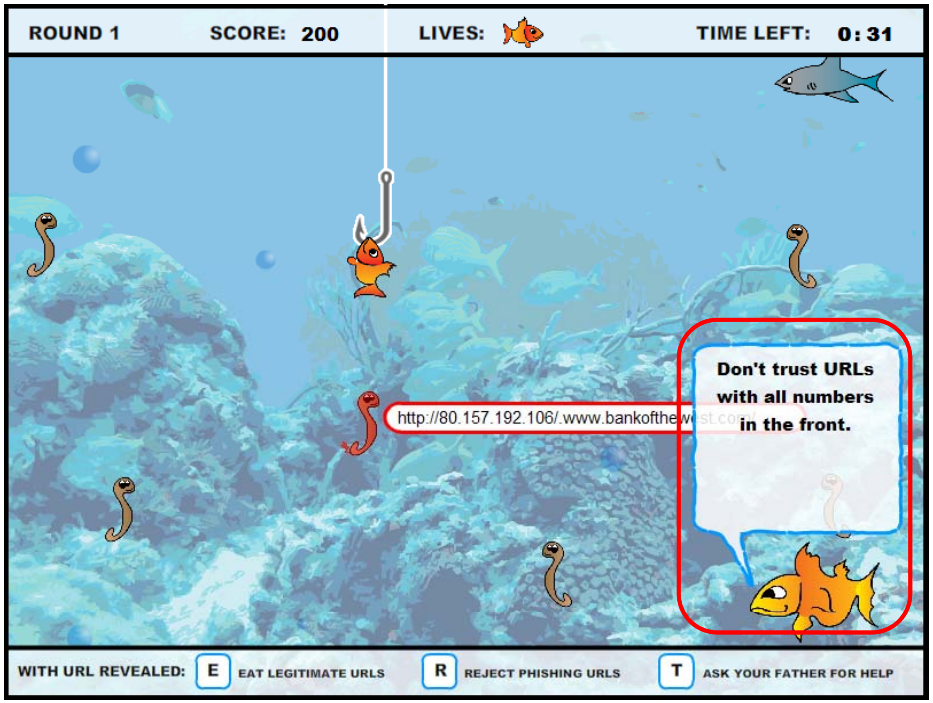
ROUND 1
SCORE: 200
LIVES:
TIME LEFT: 0:31
http://80.157.192.106/.www.bankofthe
Don't trust URLs with all numbers in the front.
WITH URL REVEALED:
E EAT LEGITIMATE URLS
R REJECT PHISHING URLS
T ASK YOUR FATHER FOR HELP

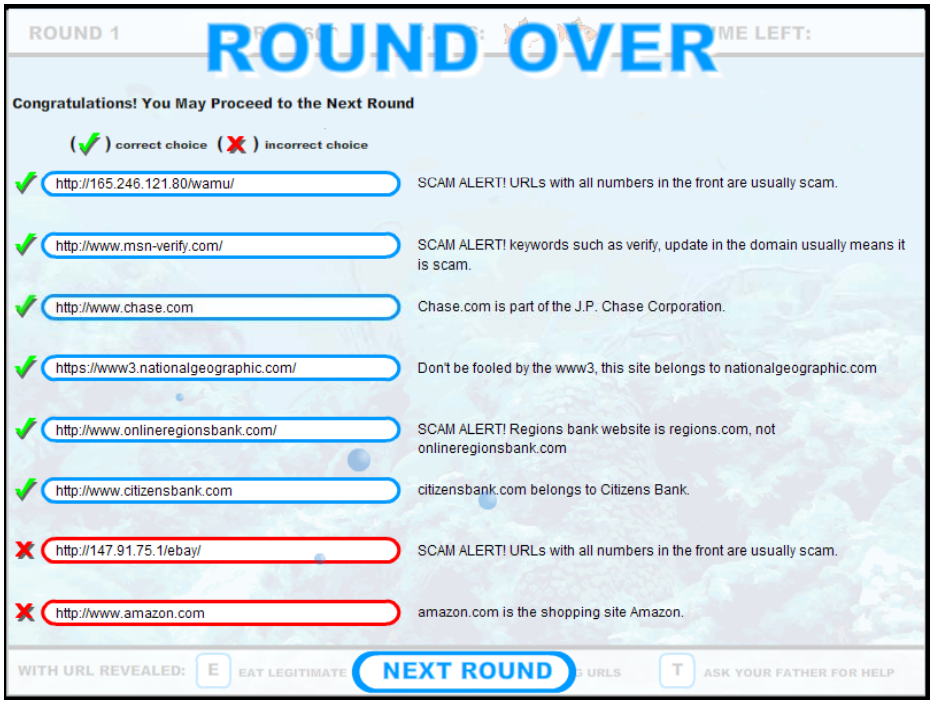
ROUND OVER
Congratulations! You May Proceed to the Next Round
(✓) correct choice (✗) incorrect choice
✓ http://165.246.121.80/wamu/ SCAM ALERT! URLs with all numbers in the front are usually scam.
✓ http://www.msn-verify.com/ SCAM ALERT! keywords such as verify, update in the domain usually means it is scam.
✓ http://www.chase.com Chase.com is part of the J.P. Chase Corporation.
✓ https://www3.nationalgeographic.com/ Don't be fooled by the www3, this site belongs to nationalgeographic.com
✓ http://www.onlineregionsbank.com/ SCAM ALERT! Regions bank website is regions.com, not onlineregionsbank.com
✓ http://www.citizensbank.com citizensbank.com belongs to Citizens Bank.
✗ http://147.91.75.1/ebay/ SCAM ALERT! URLs with all numbers in the front are usually scam.
✗ http://www.amazon.com amazon.com is the shopping site Amazon.
NEXT ROUND

More about the game

- Four rounds
  - Increasing difficulty
  - Two minutes in each round
- Eight URL “worms” in each round
  - Four phishing and four legitimate URLs
  - Users must correctly identify 6 out of 8 URLs to advance

User Study

- **Test participants’ ability to identify phishing web sites before and after training**
  - 10 URLs before training, 10 after, randomized
  - Up to 15 minutes of training
- **Three conditions**:
  - Web-based phishing education
  - Tutorial
  - Game
- **14 participants in each condition**
  - Screened out security experts
  - Younger, college students

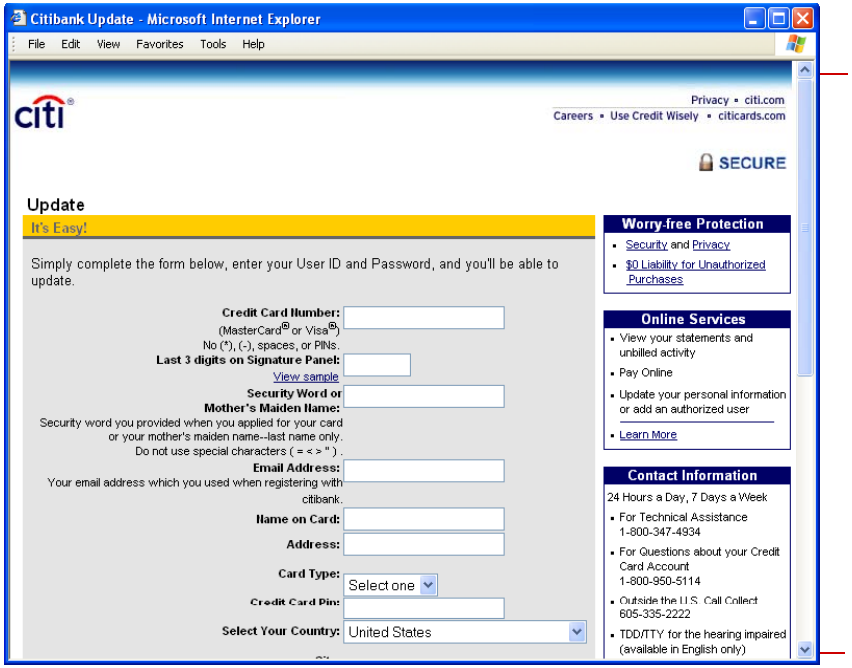

Falling for Phishing

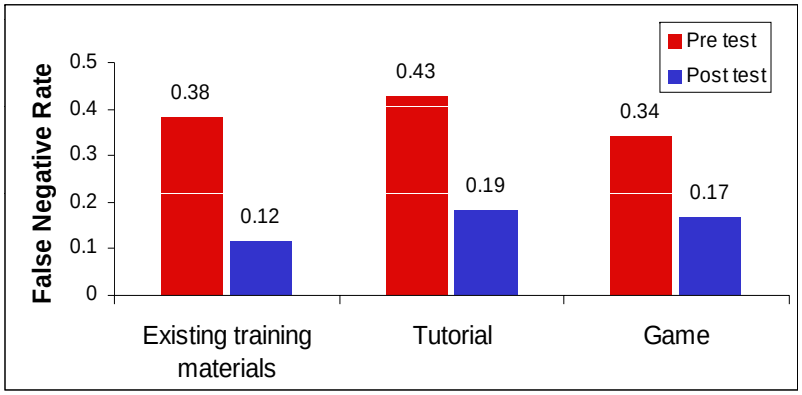

Everyone becomes more cautious

Misidentifying Legitimate Sites

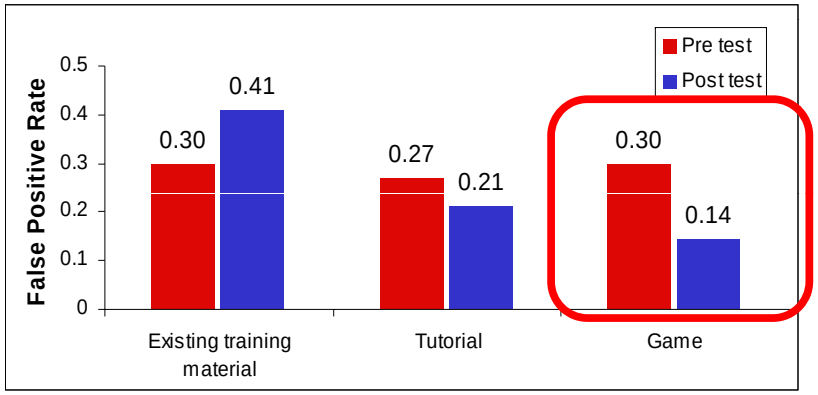

….but it does not mean that they discriminate better

Results

- No significant difference in false negatives among the three groups
- **Game group performed best in false positives**
- **Game condition performed best in total correctness**
- The training material made people paranoid but not more effective!

PhishGuru Embedded Training

- **Rationale**: Existing training doesn't work
- **Solution**: Tool to **train people during their normal use of email**
  - Periodically, insert fake phishing emails in people's regular mail
  - If person falls for it, **intervention warns and trains user** in succinct and engaging format
  - Provide reports to help companies **assess their preparedness levels**
  - Can be customized for individual orgs.

P. Kumaraguru, Y. Rhee, S. Sheng, S. Hasan, A. Acquisti, L. Cranor and J. Hong. Getting Users to Pay Attention to Anti-Phishing Education: Evaluation of Retention and Transfer. Proceedings of the 2nd Annual eCrime Researchers Summit, October 4-5, 2007

Comic Strip Intervention

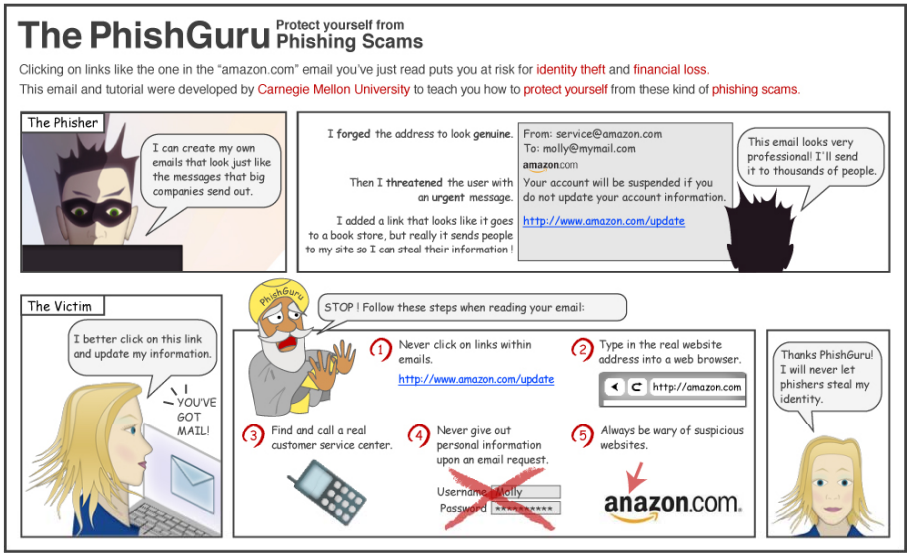

PhishGuru: Evaluation

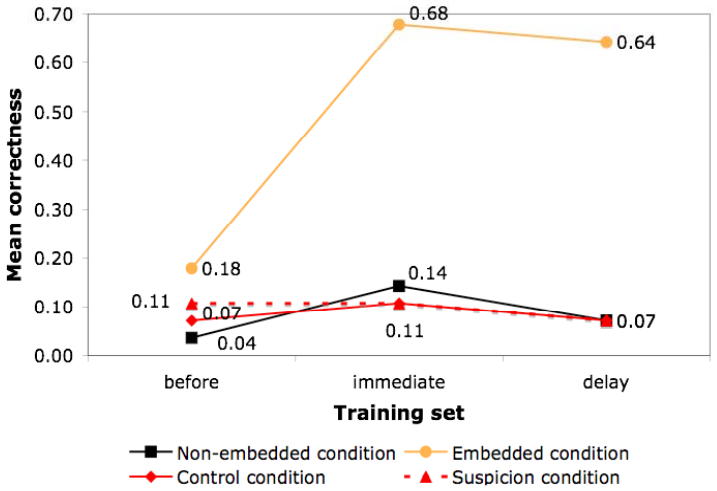

Status Today

Anti-Phishing Phil Game

- Online interactive training game
- Played by over **60000** people
- Licensed to DTCC, PT, etc.
- Versions in multiple languages
- Many inquiries

PhishGuru Training

- Embedded training
- Send fake phishing emails and train those who fall for them
- **Pilot** at Portugal Telecom
- Pilots with AT&T, CHLA, etc

PILFER Email Filter

- Outperforms top spam filters (catches more phish & 25x fewer false alarms)
- SpamAssassin plugin (potential **100mil** users)
- Pilots w. PT and CMU

CANTINA Web Filter

- Faster at detecting phish than blacklists
- More effective against spearphishing than blacklists
- **Pilot** with Portugal Telecom

Concluding Remarks

- Social engineering is not a new phenomenon
- Unfortunately, with the Web, it can be carried

Phishing: A never ending arms race?

- Many large organizations (private and government) are extremely concerned
- As technology evolves and as users become

Probably

Q&A

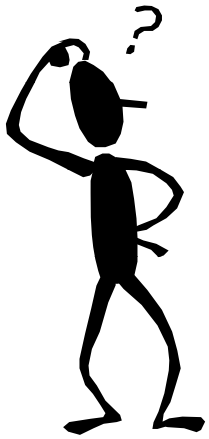

Acknowledgement: A number of the slides used in this presentation have been adapted from presentations developed by or with my collaborators in this project. This includes Lorrie Cranor, Jason Hong, Julie Downs,Ponnurangam Kumaraguru, Ian Fette and Steve Sheng.

References

- **PILFER** - Ian Fette , Norman Sadeh and Anthony Tomasic, "**Learning to Detect Phishing Emails**", Proceedings of the 16th International World Wide Web Conference, May 2007 (WWW2007): http://www.cs.cmu.edu/~sadeh/Publications/Small%20Selection/www07%20FINAL%20SUBMISSION.pdf
- **PHIL** - S. Sheng, B. Magnien, P. Kumaraguru, A. Acquisti, L. Cranor, J. Hong, and E. Nunge. Anti-Phishing Phil: The Design and Evaluation of a Game That Teaches People Not to Fall for Phish. In *Proceedings of the 2007 Symposium On Usable Privacy and Security,* Pittsburgh, PA, July 18-20, 2007. http://cups.cs.cmu.edu/soups/2007/proceedings/p88_sheng.pdf
- **PHISH GURU** - P. Kumaraguru, Y. Rhee, S. Sheng, S. Hasan, A. Acquisti, L. Cranor and J. Hong. Getting Users to Pay Attention to Anti-Phishing Education: Evaluation of Retention and Transfer. Proceedings of the 2nd Annual eCrime Researchers Summit, October 4-5, 2007, Pittsburgh, PA, p. 70-81. http://www.ecrimeresearch.org/2007/proceedings/p70_kumaraguru.pdf
- **CANTINA** - Y. Zhang, J. Hong, and L. Cranor. CANTINA: A content-based approach to detecting phishing web sites. In *Proceedings of the 16th International conference on World Wide Web,* Banff, Alberta, Canada, May 8-12, 2007.

Press Coverage

Web tool detects something phishy by Bonnie Pfister, *Pittsburgh Tribune-Review*, 12/11/07.

Researchers: Current education inadequate to fight phishing by Elizabeth Montalbano, *Computerworld,* 10 October 2007.

Online game teaches users about the threats of phishing by Kun Xian Leong, *The Tartan,* 8 October 2007.

Phishing victims learn online security lesson by Robert Jacques, *vnunet.com,* 3 October 2007.

Fighting Phish, *Dr. Dobb's Journal,* 2 October 2007.

Coolest Security Tool Ever! - Online game to teach cyber-security by Alexandru Dumitru, *Softpedia,* 2 October 2007.

Anti-Phishing Game To Help Raise Awarness, *Portalit.net,* 1 October 2007.

Scientists develop Anti-Phishing game to educate Web users by Ruben Francia, *BLORGE.com,* 29 September 2007.

Carnegie Mellon's Online Game Helps People Recognize Internet Scams, Phishing by Regina Sass, *Associated Press,* 28 September 2007.

A new game developed at Carnegie Mellon University educates users on phishing threats. by Christopher Nickson, *Digital Trends,* 27 September 2007.

Phishers caught hook, line and sinker by Stuart Turton, *PC Pro,* 26 September 2007.

Carnegie Mellon floats anti-phishing game by Robert Jaques, *PC Magazine,* 26 September 2007.

CMU's Anti-Phishing Phil helps users identify Internet scams--try it! by Deb Smit, *POP City,* 26 September 2007.

Fish named Phil helps foil phishers, *CBC News,* 26 September 2007.

The truth about anti-phishing toolbars by Paul Roberts, *InforWorld Tech Watch,* 30 November 2006.

Study blasts failing phishing toolbars by Shaun Nichols, *vnunet.com,* 22 November 2006.

Phishing toolbars: all as hopeless as one another by John E. Dunn, *Techworld,* 20 November 2006.

Phishing Filter Prevents E-Mail Identity Theft by Brian Livingston, *Executive Tech,* 18 July 2006.

Don't click anything! by Thomas Olson, *Pittsburgh Tribune-Review,* 12 July 2006.

Researchers work to thwart cleverer cyber scammers by Corilyn Shropshire, *Pittsburgh Post-Gazette,* 12 July 2006.