What we now know about phishing websites

Richard Clayton

(joint work with Tyler Moore)

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Yahoo!
19th August 2008

Academics & phishing

- Everyone can play! Display instant expertise!!
  - examine psychology, attempt to block spam, detection of websites, browser enhancements, password mangling, reputation systems etc
- Our approach : Security Economics
  - phishing will continue, because humans involved!
  - so we measure the impact, assess the effectiveness of countermeasures, work out how to change incentives so that problem tends to fix itself…

Academics & the real world

- Papers have to be “novel research”
- PhDs have to be “a contribution”
- Where the “real world” is tackled, tendency to pick the “low hanging fruit” and move on
- “Peer review” process requires peers
- Natural tendency not to want to report failures
- Natural tendency not to admit mistakes

Last year's "Summary"

- Take-down has an impact
  - but it is not fast enough to make losses zero
- Rock-phish gang have a good recipe
  - planned ? or just stumbled upon ?
- Wide variations in bank performance
  - incompetence? or facing better attackers?
- Some "phishing losses" are indeed phishing
  - but sums too rough to discount key-loggers &c

Data Sources

- Originally mining PhishTank dataset
  - free and apparently accurate and substantial
- Now getting data from a brand owner and two brand protection companies (plus PhishTank and “Artists Against 419”)
- These phishing “feeds” have common components but turn out to be different…

	PhishTank	BrandProtectA
URLs	10924	13318
Non-duplicate URLs	8296	8730
Unique URLs	3019	2585
Rock-phish domains	586	1003
Unique rock-phish domains	127	544

63% of total overlap (9380 URLs) from “PhishReporter”

remainder from 316 separate submitters

Verification time (average)	46 hours	8 seconds
Verification time (median)	15 hours	8 seconds

PhishTank errors (July/Aug 07)

- Errors in submissions: (44% from single submitters, but 1.2% from most active)
- Errors in voting: 39 false +ve, 3 false –ve
- Inaccuracy of voting (count disagreements):
  - fewer than 100 votes: 14% of time
  - most active voters: 3.7% of time
- “High-conflict” users make the same mistakes

Attacks on PhishTank

#1 Submit invalid reports

#2 Vote for phish as not-phish

#3 Vote for not-phish as phish

- Hard to defend
  - power-law distribution of submissions, and also in voting participation
  - easy to get an accurate reputation (97% phish)
  - failure to canonicalise rock-phish

“Wisdom of Crowds” & security

- Distribution of user participation matters
  - power laws puts power into hands of the few
  - however, you do want keen people...
- Decisions must be difficult to guess
  - you want people participating not robots
- Do not make users work harder than needed
  - canonicalise the data

Feeds are not shared

- Brand-protection companies obtain feeds from many places (including PhishTank)
- They run their own detectors
- They sell feeds, but don’t share them
- Hence Company A, who sells services to Bank A1, can be unaware of sites detected by Company B – and doesn’t take them down

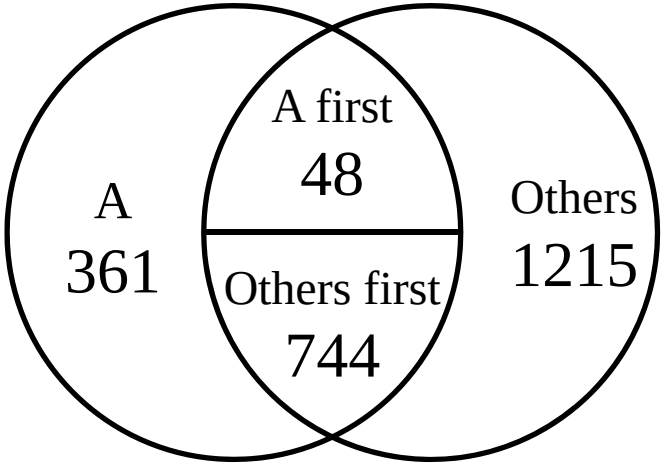

Ordinary phishing sites

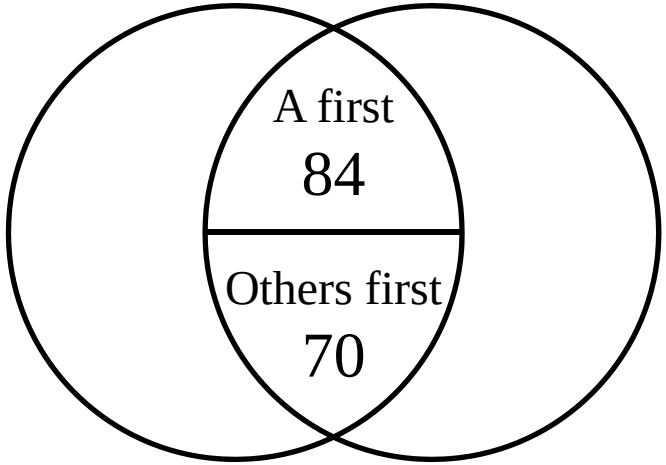

Delay in detecting (hours)

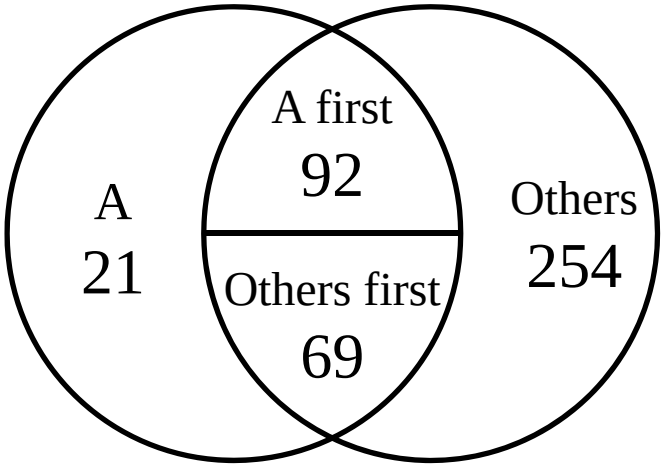

Mean lifetime (hours)

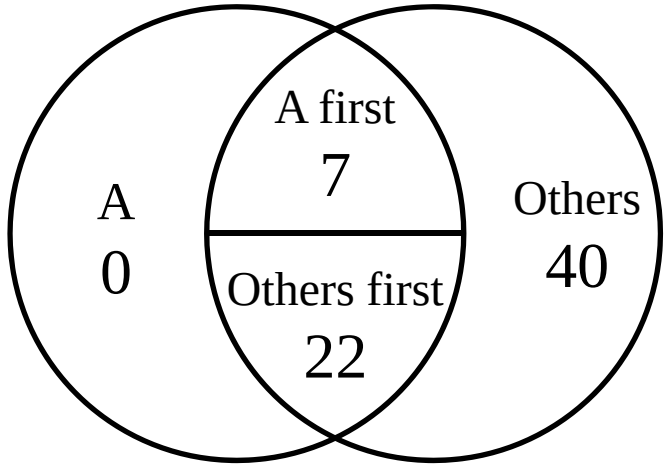

Median lifetime (hours)

Bank A1’s experience as a client of BrandProtection company A

Company A v Company B

- Same pattern continues for top 6 banks for Company A and B, and for all *n* clients
- However, less pronounced for B: which seems to have a better feed [or maybe just one that is much more aligned with ours!]
- But A's clients bigger and proportion missed goes up with size; so B's prowess may be more a structural issue than just extra effectiveness

Thi$ repreent ri$k

- Longer lifetimes => more visitors (Webalizer logs)
- Hence we can assess impact of longer lifetimes:

Exposure figures (6 month totals)	A's banks		B's banks	
	Khour	$m	Khour	$m
Actual values	1005	276	78	32
Expected if sharing	418	113	61	28.5
Effect of no sharing	**587**	**163**	**17**	**3.5**

Hence…

- Banks should force brand-protection companies to share feeds
  - cf the anti-virus community for last 15 years
- Brand-protection companies could form a "club" to prevent new entrants from free-riding
  - don't have to make feeds free, just share them
- Side-note: free-riding by rock-phish attacked banks only works some of the time!

Types of phishing website

- Insecure end user or machine (76% of sites)
  - **`http://www.example.com/~user/www.bankname.com/`**
  - **`http://www.example.com/bankname/login/`**
- Free web hosting (17% of sites)
  - **`http://www.bank.com.freespacesitename.com/`**
- Misleading domain name (unusual)
  - **`http://www.banckname.com/`**
  - **`http://www.bankname.xtrasecuresite.com/`**
- Random domains (after canonicalisation)
  - rock-phish 4%, fast-flux 1.4%, “ark” 1.4%

How are insecure machines found?

- Traditionally machines found by “scanning” hence interest in Intrusion Detection Systems, “slow scan” software etc etc
- We have been collecting Webalizer logs (wanted to count number of visitors to sites and hence calculate impact of prompt take-down)
- Webalizer parses referrer strings to determine search terms used to locate the sites….

Typical searches in weblogs

- Hand categorisation, but most were obvious
  - many searches for MP3s ! these were ignored
- Vulnerability
  - **`phpizabi v0.848b c1 hfp1`** (CVE-2008-0805)
- Compromise
  - **`allintitle:welcome paypal`**
- Shell
  - **`c99shell drwxrwx`**

Webalizer logs (June 07 – March 08)

- 2486 domains with world-readable logs
- 1320 (53%) had one or more search terms
- 25 cases where searches provably linked

	Domains	Phrases	Visits
Any evil search	204	456	1207
Vulnerability search	126	206	582
Compromise search	56	99	265
Shell search	47	151	360

More statistics

- Assume Webalizer sites are a random sample of all sites (make up your own mind on that)
  - if so, then 95% confidence interval for incidence of "evil searching" (aka "dorks") is 15.3% to 19.8%
- Did our own searches (thanks Yahoo!) on evil and non-evil terms and checked if phishing site
  - 1.9% sites found with evil terms used for phishing
  - 0.73% sites with non-evil terms (statistically significant difference)

Overlap of search results

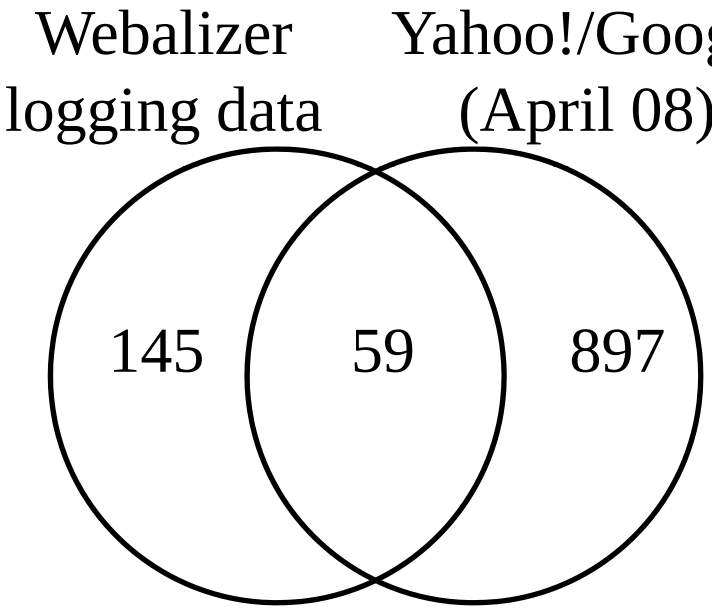

Many searches don’t work any more, but lots more sites to attack!

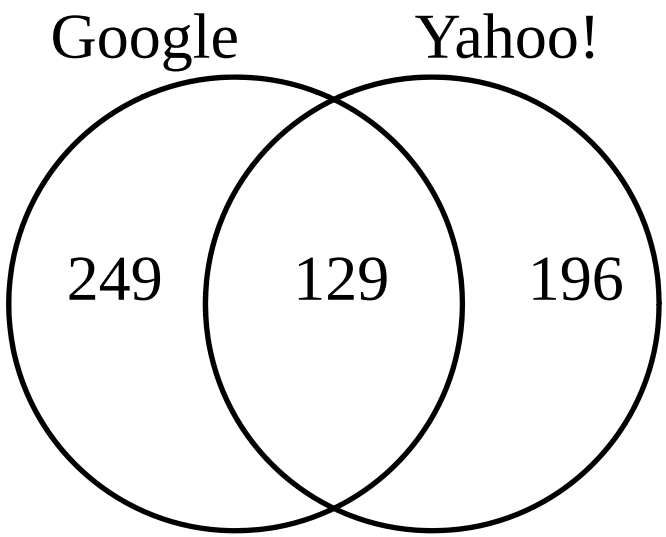

There’s a surprising lack of overlap in the results

Recompromise

- Consider phishing pages on same site more than a week apart (likely a different attacker)
- 9% of all sites recompromised within 4 weeks, rising to 19% within 24 weeks
- For Webalizer sites this is 15% rising to 33%
- If evil search terms present then this becomes 19% rising to 48% (14% to 29% if no terms)
- This doubling is statistically significant!

Comparing take-down times

- Defamation – believed to be quick (days)
- Copyright violation – also prompt(ish)
  - experimentally “days” (albeit with prompting)
- Fake escrow agents
  - average 9 days, median 1 day
  - note that AA419 aware of around 25% of sites
- Mule recruitment sites (Sydney Car Center etc)
  - average 13 days, median 8 days

Phishing Lifetimes (hrs)	sites	mean	median
Free-web hosting			
all	395	47.6	0
brand-owner aware	240	4.3	0
brand-owner unaware	155	114.7	29
Compromised machines			
all	193	49.2	0
brand-owner aware	105	3.5	0
brand-owner unaware	155	103.8	10
Rock-phish domains	821	70.3	33
Fast-flux domains	315	96.1	25.5

Incentives

- Most of the take-down time variations are explainable in terms of incentives
  - the motivated complain again&again until removed
  - the banks are ignoring mule recruitment (not their problem) so just volunteers (vigilantes)
  - escrow faster than mule sites: attacking the innocent?
    or maybe escrow.com is doing more than we think?
  - no-one's job to remove fake pharmacies (and no active volunteers) so their lifetime is ~2 months

Child Sexual Abuse Images (“CAI”)

- Provided with anonymised data by IWF
- Jan–Dec 2007 2585 domains
  - ignoring 8 (free-web?) domains with >100 reports
- Computed the initial take-down time (ignored recompromise): mean 21 days, median 11 days
- If we include sites with no removal at all then mean grows to 30 days (and counting)
  - median also grows by one day

Why so slow?

- In fact quick within the UK : IWF checks with police and then contacts the ISP
- But “not authorised” to act internationally
- Passes data via UK police to foreign forces
  - but may not reach local field office for a while
- Also pass to another INHOPE member
  - but (eg) NCMEC only act “when appropriate”
- Confusion of aims (removal/catch criminals)

Ongoing research agenda

- How many phishers are there ?
- How much phishing is phishing ?
- How do we fix the incentives to prevent phishing from being effective ?
- Phishing is now mechanised and uses standard kits – we'd like to disrupt them!
- Phishing attacks also involve spam: the timing of this is as relevant as site take-down times

2008 summary

- “Wisdom of crowds” is not a security panacea
- The phishing site take-down industry is putting significant funds at risk by not co-operating
- Search engines are widely used to find websites to compromise (and re-compromise)
- Takedown times are affected more by incentives than by formal structures
- Slowness of removal of CAI is a scandal

What we now know about phishing websites

BLOG: http://www.lightbluetouchpaper.org/

http://www.cl.cam.ac.uk/~rnc1/
http://people.seas.harvard.edu/~tmoore/

PAPERS: http://www.cl.cam.ac.uk/~rnc1/publications.html

UNIVERSITY OF CAMBRIDGE
Computer Laboratory