Global Phishing Survey: Domain Name Use and Trends in 1H2008

An
APWG
Industry
Advisory

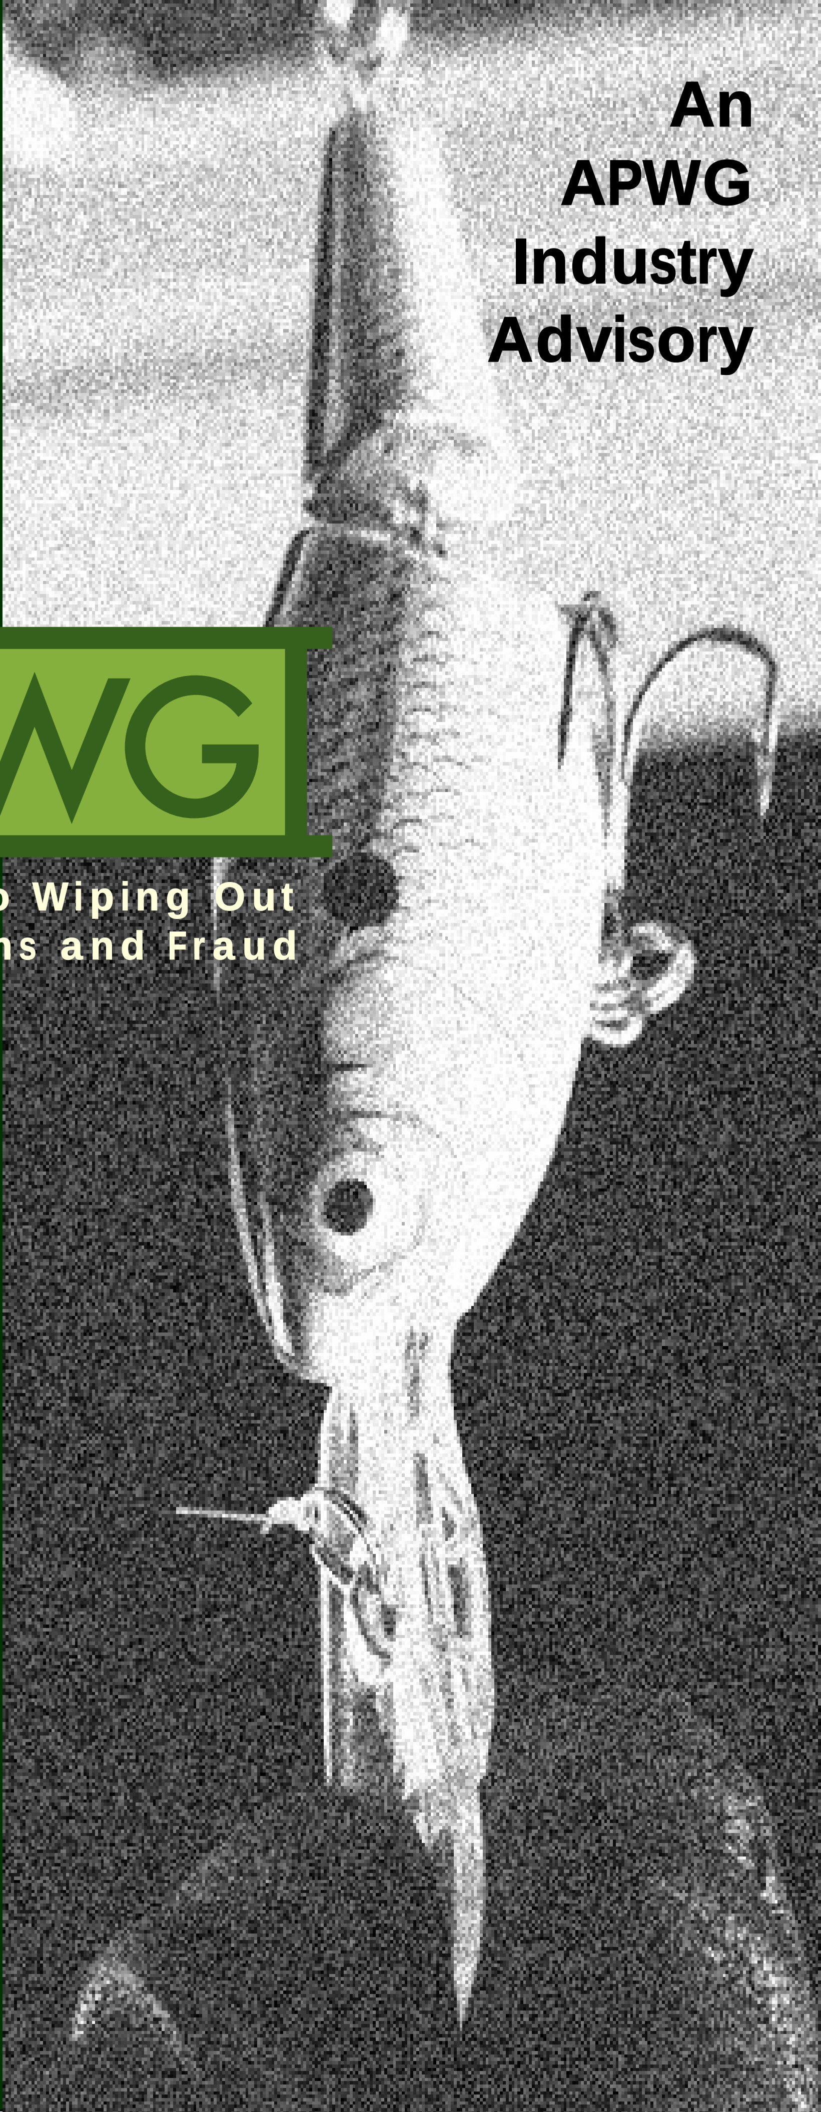

APWG

Committed to Wiping Out
Internet Scams and Fraud

November 2008

Authors:

Rod Rasmussen
Internet Identity
<rod.rasmussen at internetidentity.com>

Greg Aaron
Afilias
<gaaron at afilias.info>

Table of Contents

Disclaimer: PLEASE NOTE: The APWG and its cooperating investigators, researchers, and service providers have provided this study as a public service, based upon aggregated professional experience and personal opinion. We offer no warranty as to the completeness, accuracy, or pertinence of these data and recommendations with respect to any particular company's operations, or with respect to any particular form of criminal attack. Please see the APWG website – apwg.org - for more information.

Overview

Phishers are constantly experimenting and adapting. In order to combat them effectively, it is important to understand how they are using domain names and why. Domain name usage is an important measure of the scope of the global phishing problem, and examination of domain name trends can provide new anti-abuse strategies.

This study describes our analysis of a comprehensive database of phishing that took place in the first half of 2008 (1H2008), and is a follow-up to our 2007 study.[1] Specifically, the data in this new report includes all the phishing attacks detected between January 1, 2008 and June 30, 2008, as collected by the APWG and supplemented with additional reports from several phishing feeds and private sources. The APWG phishing repository is the Internet's most comprehensive archive of e-mail fraud and phishing activity.[2]

New to this 1H2008 report are attack statistics, and measurements of phishing site up-times. Our data reveals some new trends, and we hope that bringing these tactics to light will lead to improved anti-phishing measures.

Our major findings are:

1. Phishers continue to target specific Top-Level Domains (TLDs) and specific domain name registrars, and shift their preferences over time. Metrics that measure the pervasiveness of phishing in TLDs provide a valuable way to identify exploitation by phishers who register domain names.
2. Anti-phishing programs implemented by domain name registries can have a noticeable effect on the up-times (durations) of phishing attacks. We see some direct correlation between the efforts of several large gTLD and ccTLD operators and the amount of time that phishing sites remained live within their TLDs.
3. Phishers are engaged in the large-scale use of subdomain services to host and manage their phishing sites. Such attacks even account for the majority of attacks in certain large TLDs.

Basic Statistics

Millions of phishing URLs were reported in 1H2008, but the number of phishing attacks and domain names used to host them is much smaller. This is due to several factors:

[1] http://www.antiphishing.org/reports/APWG_GlobalPhishingSurvey2007.pdf
[2] Our study is designed to complement rather than duplicate the APWG's quarterly Phishing Activity Trend reports, which measure metrics including the number of unique phishing reports received per month, the number of brands attacked per month, and the countries where phishing sites were hosted. These reports are available at: http://www.apwg.org/phishReportsArchive.html

1. Some phishing involves customized attacks by incorporating unique numbers in the URLs (to track targeted victims, or to defeat spam filters). A single phishing attack can therefore be seen as thousands of individual URLs.
2. Phishers often use one domain name to host simultaneous attacks against multiple target brands. For example, the Rock Phish gang is known for placing four or more different phishing attacks on each domain name it registers.
3. A phishing site may have multiple pages, each of which may be reported.

The 1H2008 data set yielded the following basic statistics:

- There were at least **47,324 phishing attacks**. An "attack" is defined as a phishing site that targets a specific brand or entity. A domain name can host several discrete attacks against different banks, for example.
- The attacks occurred on **26,678 unique domain names**.[1] This is slightly down from 2H2007, when 28,818 domain names were used.
- In addition, phish were found on **3,389 unique IP addresses**, rather than on domain names. (For example: http://91.121.81.84/do.php?cmd=SignIn.) This is down significantly from 2H2007, when 5,217 unique IP addresses were used, and down from 1H2007, when 6,336 unique IPs were used.
- Phishing took place on domain names in **155 TLDs**. This is up from 2H2007, when only 145 TLDs were used.
- Only 52 of the 29,073 domain names were Internationalized Domain Names (IDNs). These mostly involved .HK domain names used by the Rock Phish gang early in 2008.

Overall Statistics: 1H2008 versus 2H2007

	1H2008	2H2007
Phishing domain names:	26,678	28,818
IP-based phish (unique IPs):	3,389	5,217
TLDs phished on:	155	145
"Attacks":	>47,342	
IDN domains:	52	10

Each domain name's registrar of record was usually not reported at the time of the phish. In most registries, a domain name can have multiple "lifetimes" as the name is registered, is deleted or expires, and is then registered anew. Obtaining accurate registrar sponsorship of a domain name requires either time-of-attack WHOIS data, or historical registry-level data. This data has not been collected in a methodical or comprehensive manner by the anti-phishing community. Registrar-specific statistics and trends are certainly of interest, and are an opportunity for future studies.

[1] "Domain names" are defined as second-level domain names, plus third-level domain names if the relevant registry offers third-level registrations. An example is the .CN (China) registry, which offers both second-level registrations and third-level registrations (in zones such as com.cn, gov.cn, zj.cn, etc.). However, see the "Subdomains Used for Phishing" section for commentary about how these figures may undercount the phishing activity in a TLD.

Prevalence of Phishing by Top-Level Domain (TLD)

We analyzed the 26,678 phishing domains to see how many fell into which TLDs. The absolute counts by TLD are interesting, but the sizes of the various TLDs vary widely. To place the numbers in context and measure the *prevalence* of phishing in a TLD, we use the metrics "Phishing Domains per 10,000" and "Phishing Attacks per 10,000."

"Phishing Domains per 10,000" is a ratio of the number of domain names used for phishing in a TLD to the number of registered domain names in that TLD.[1] This metric is a way of revealing whether a TLD has a higher or lower incidence of phishing relative to others. In 1H2008, phishing occurred on domain names in 155 TLDs. Of these, we were able to obtain the domain count statistics for 109 TLD registries.[2] Those 109 TLDs contained 98% of the phishing domains in our data set (26,026 out of the 26,678), and a total of 167,638,848 domain names overall. Industry estimates put the total number of domain names in existence worldwide at the end of 1H2008 at approximately 168,000,000.[3]

The complete tables are presented in Appendix A, including the scores and the number of phish in each TLD.

- The **median score was 2.3**.
- The **average score was 9.21**, which was skewed by a few high-scoring TLDs.
- .COM, the world's largest and most ubiquitous TLD, had a score of 1.6. .COM contains 45.9% of the phishing domains in our data set, and 45.7% of the domains in the TLDs for which we have domains-in-registry statistics. In the ranking of TLDs by score, there are 34,536,256 domains in the TLDs ranked below .COM, and 56,442,122 in the TLDs ranked above .COM.

We therefore suggest that *scores between .COM's 1.6 and the median of 2.3 occupy the middle or "normal" ground*, with scores above 2.3 indicating TLDs with increasingly prevalent phishing.

The metric "Phishing Attacks per 10,000" provides insight into what TLDs are predominantly used by phishers who use subdomain services, and where phishers (notably the Rock Phish gang) place multiple phish on one domain. New Rock Phish attacks were evident in .BE and .UK, for example, and declining Rock Phish attacks were clearly seen in .HK.

Notes regarding the statistics:

- What explains why a TLD has a higher or lower phishing score, and what do the scores mean for registry operators and anti-phishing efforts? For more background on factors that can affect a TLD's score, please see our 2007 study, at: http://www.antiphishing.org/reports/APWG_GlobalPhishingSurvey2007.pdf

[1] Score = (phishing domains / domains in TLD) x 10,000
[2] For the purposes of this study, we used the number of domain names in each registry in May 2008. Sources: ICANN.org (for gTLD and sTLD monthly registry reports), ccTLD registry operators, Latinoamericann.org, Webhosting.info.
[3] VeriSign and Zooknic, http://www.verisign.com/static/044191.pdf

- A small number of phish can increase a small TLD's score significantly, and these pushed up the study's median score. The larger the TLD, the less a phish influences its score, and indeed the largest TLDs tend to appear lower in the rankings.
- A registry's score can be increased by the action of just one phisher, or one vulnerable or inattentive registrar.

Eliminating TLDs that had less than 30,000 domains under management or less than 25 phishing domains yields the following:

Top 20 Phishing TLDs by Score

Minimum 25 phishing domains and 30,000 domain names in registry

Rank	TLD	TLD Location	# Unique Phishing attacks 1H2008	Unique Domain Names used for phishing 1H2008	Domains in registry in May 2008	Score: Phish per 10,000 domains 1H2008
1	hk	Hong Kong	2,278	516	160,336	**32.2**
2	th	Thailand	154	84	35,757	**23.5**
3	bz	Belize	52	43	43,216	**10.0**
4	ve	Venezuela	86	71	75,000	**9.5**
5	cl	Chile	274	128	212,153	**6.0**
6	ro	Romania	184	142	284,700	**5.0**
7	li	Liechtenstein	97	26	59,546	**4.4**
8	name	sponsored TLD	331	126	289,343	**4.4**
9	tw	Taiwan	319	145	385,500	**3.8**
10	kr	Korea	697	345	945,000	**3.7**
11	es	Spain	883	333	970,580	**3.4**
12	in	India	252	150	454,330	**3.3**
13	mx	Mexico	122	83	255,406	**3.2**
14	sk	Slovakia	63	50	159,758	**3.1**
15	pl	Poland	417	300	960,000	**3.1**
16	gr	Greece	82	60	203,000	**3.0**
17	ru	Russia	1,907	362	1,427,928	**2.5**
18	hu	Hungary	120	85	372,700	**2.3**
19	org	generic TLD	2,384	1,425	6,905,011	**2.1**
20	net	generic TLD	4,159	2,305	11,623,856	**2.0**

The "generic" TLDs are used by and are popular with registrants across the world. There is some variance in their scores:

Phishing in gTLDs by Score

Rank	TLD	Domains in registry, May 2008	Domain names used for phishing, 1H2008	Score: Phishing domains per 10,000
57	.org	6,905,011	1,425	**2.1**
59	.net	11,623,856	2,305	**2.0**
66	.biz	2,035,357	353	**1.7**
71	.com	76,625,770	12,275	**1.6**
80	.info	5,042,032	684	**1.4**

If measured by Attack Score, certain TLDs vault into much higher rankings:

Top 15 Phishing TLDs by Attack Score

Minimum 50 phishing attacks and 30,000 domain names in registry

Rank	TLD	TLD Location	# Unique Phishing attacks 1H2008	Unique Domain Names used for phishing 1H2008	Domains in registry in May 2008	Score: Phish per 10,000 domains 1H2008	Score: Attacks per 10,000 domains 1H2008
1	hk	Hong Kong	2,278	516	160,336	32.2	**142.1**
2	th	Thailand	154	84	35,757	23.5	**43.1**
3	su	Soviet Union	154	14	60,543	2.3	**25.4**
4	li	Liechtenstein	97	26	59,546	4.4	**16.3**
5	ru	Russia	1,907	362	1,427,928	2.5	**13.4**
6	cl	Chile	274	128	212,153	6.0	**12.9**
7	bz	Belize	52	43	43,216	10.0	**12.0**
8	ve	Venezuela	86	71	75,000	9.5	**11.5**
9	name	sponsored TLD	331	126	289,343	4.4	**11.4**
10	fr	France	1,236	107	1,128,776	0.9	**10.9**
11	es	Spain	883	333	970,580	3.4	**9.1**
12	be	Belgium	690	62	791,737	0.8	**8.7**
13	tw	Taiwan	319	145	385,500	3.8	**8.3**
14	kr	Korea	697	345	945,000	3.7	**7.4**
15	ro	Romania	184	142	284,700	5.0	**6.5**

.SU, .RU, and .FR received high Attack Scores because phishers launched large numbers of attacks in these TLDs via subdomain hosting services. (For more, see “Use of Subdomains for Phishing,” below.) .BE had an elevated Attack Score because .BE domains were used extensively by the Rock gang, which placed multiple phish on many of the domain names it registered. Attack Score is therefore a useful measure of the pervasiveness of phishing in a namespace.

The .SU domain is notable because it was to have been phased out years ago, after the dissolution of the Soviet Union. However, it remains in the DNS root and is accepting new registrations. .SU is managed by the Russian Institute for Public Networks, which also operates the .RU TLD.[1]

.UK was targeted by the Rock Phish gang, which exploited at least one slow-to-respond registrar.

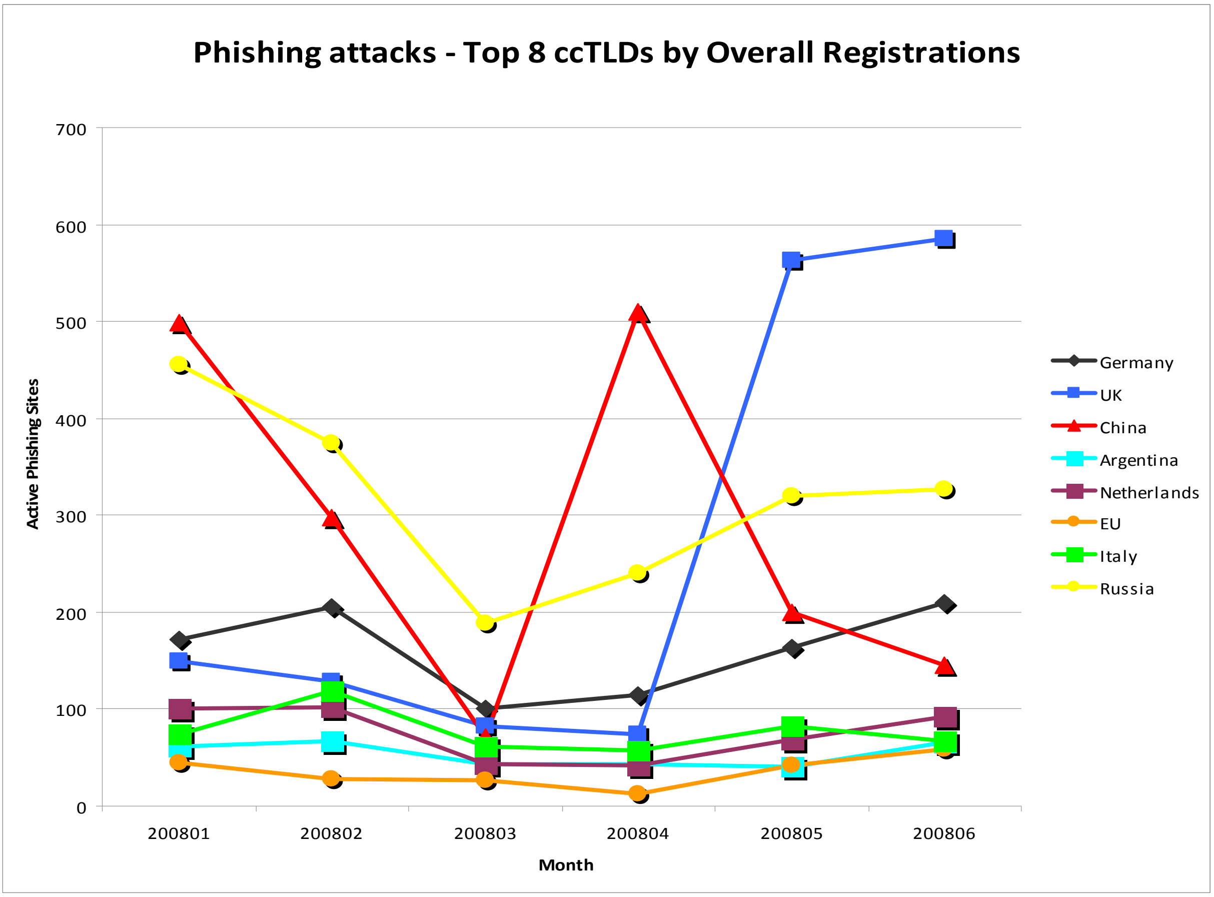

High-scoring TLDs almost invariably suffered from the systematic exploitation by phishers, and highlight how a single point of vulnerability can lead to significant problems. A few examples are:

- **.HK** (Hong Kong. Score 32.2; 2,278 attacks; 516 phishing domains out of 160,336 in the registry.) .HK was targeted by the Rock Phish gang beginning in 2007, and the .HK registry's efforts to defeat those attacks bore fruit early in 2008.[2],[1]

[1] http://en.wikipedia.org/wiki/.su

[2] https://www.hkdnr.hk/company_info/pressrelease.jsp?item=40

- **.TH** (Thailand. Score: 23.5; 154 attacks; 84 phishing domains out of 35,757 domains in the registry). Forty-six of the the phishing domains were under the AC.TH (academic) zone, and 23 more were under the GO.TH (government) zone. In other words, phishers were systematically taking advantage of insecure institutional servers rather than fraudulently registering new .TH domains for their attacks.
- **.LY** (Libya. Score 122.6; 39 attacks; 39 phishing domains out of 3,181 domains in the registry.) The phishing domains were maliciously registered in the BIZ.LY zone, and targeted one brand.

How long did the phishing attacks last, and how damaging were they? To learn more, we next analyzed uptimes.

Phishing By Uptime

In 1H2008, Internet Identity monitored the "uptimes" or "live" times of the phishing attacks in the data set. Uptimes are a vital measure of how damaging phishing attacks are, and are a measure of the success of mitigation efforts. The raw number of phishing attacks is an important measure, but the mitigation of those attacks (or lack thereof) is the key factor for both the criminals and their victims. The longer a phishing attack remains active, the more losses that accrue. For example, a top-ten American bank estimates that *at least US$300 is lost for every hour that a phishing site remains up*.[2]

Phishers therefore strive for maximum uptime, and choose domain spaces and providers accordingly. Phishers prefer vulnerable or inattentive registrars and registries, and some phishers use fast-flux hosting to extend uptimes. Phishing hosted on fast-flux networks often stay up at least twice as long as those on conventional hosting. [3] Long-lived phish can skew the averages considerably, as some phishing sites may last weeks or even months. Thus medians may be a useful barometer of overall mitigation efforts.

The system used to track the uptimes automatically monitored the phishing sites, and monitoring began as soon as the system became aware of a phish via feeds or honeypots. Each phish was checked several times per hour to confirm its availability, and was not declared "down" until it has stayed down for at least one hour. (This requirement was used because some phish, especially those hosted on botnets, may not resolve on every attempt but in general remain live.) This estimate tends to under-count the "real" uptime of a phishing site, since more than 10% of sites "re-activate" after one hour of being down. However, our method is a consistent measure that allows direct comparison across

[1] http://www.antiphishing.org/reports/APWG_GlobalPhishingSurvey2007.pdf

[2] This estimate posits that the average loss from a stolen bank access credential (either online account access, a debit card, or credit card) is US$400, and that the phisher steals two such valid credentials every three hours. This impact generally holds throughout the first 72 hours of phishing site uptime. Note these are conservative estimates since they measure only are bottom-line losses, and do not factor in "soft costs" like customer support calls, unseen losses through untracked channels, or the impact of ID theft upon the customer.

[3] For excellent analyses of this phenomenon, see "Examining the Impact of Website Take-down on Phishing" by Tyler Moore and Richard Clayton: http://www.cl.cam.ac.uk/~rnc1/ecrime07.pdf and "As the Net Churns: Fast-Flux Botnet Observations" by Jose Nazario and Thorsten Holz: http://honeyblog.org/junkyard/paper/fastflux-malware08.pdf

incidents and should be fair for relative comparisons.

We calculated the average and median uptimes for all of the 1H2008 attacks, and also for the attacks associated with some of the larger TLDs. For all 47,342 attacks, the **average was 49.5 hours**, with a **median of 19.5 hours**.

The uptimes for all phishing in 1H2008, and for phish in large TLDs, were as follows:

All TLDs	Average Uptime (HH:MM:SS)	Median Uptime (HH:MM:SS)
Jan	54:33:10	21:14:55
Feb	37:44:43	17:29:40
Mar	38:45:45	14:34:03
Apr	50:07:46	18:23:37
May	56:39:48	19:41:19
June	53:55:46	22:39:09

.COM	Average Uptime (HH:MM:SS)	Median Uptime (HH:MM:SS)
Jan	68:06:45	24:10:46
Feb	44:23:57	21:19:14
Mar	42:17:44	14:30:19
Apr	54:53:48	20:36:55
May	64:09:44	21:51:05
June	57:50:23	23:20:35

.ORG	Average Uptime (HH:MM:SS)	Median Uptime (HH:MM:SS)
Jan	69:02:25	23:13:34
Feb	40:12:15	17:09:30
Mar	46:54:11	16:32:13
Apr	49:42:01	21:05:59
May	49:25:30	20:12:10
June	52:54:36	21:35:50

.INFO	Average Uptime (HH:MM:SS)	Median Uptime (HH:MM:SS)
Jan	25:35:05	21:11:12
Feb	20:23:01	17:40:54
Mar	16:10:50	12:23:56
Apr	32:35:14	11:59:21
May	55:11:20	23:58:14
June	29:40:30	12:35:20

.BIZ	Average Uptime (HH:MM:SS)	Median Uptime (HH:MM:SS)
Jan	30:29:43	16:48:53
Feb	46:32:46	14:21:16
Mar	34:19:14	15:44:08
Apr	28:40:23	5:25:15
May	57:04:56	13:15:53
June	44:20:13	18:08:23

.CN	Average Uptime (HH:MM:SS)	Median Uptime (HH:MM:SS)
Jan	43:18:11	16:11:33
Feb	20:18:53	8:43:36
Mar	48:12:15	22:42:59
Apr	16:33:47	10:10:07
May	38:52:54	10:20:21
June	43:35:26	17:53:37

.US	Average Uptime (HH:MM:SS)	Median Uptime (HH:MM:SS)
Jan	52:58:54	20:16:28
Feb	43:02:39	21:35:23
Mar	72:24:48	15:22:18
Apr	64:04:38	30:18:09
May	42:52:31	13:46:15
June	57:25:35	19:09:45

.UK	Average Uptime (HH:MM:SS)	Median Uptime (HH:MM:SS)
Jan	76:20:37	19:29:19
Feb	41:41:58	24:00:57
Mar	37:56:24	17:10:55
Apr	62:14:58	29:26:26
May	34:24:30	13:55:40
June	42:44:06	22:46:33

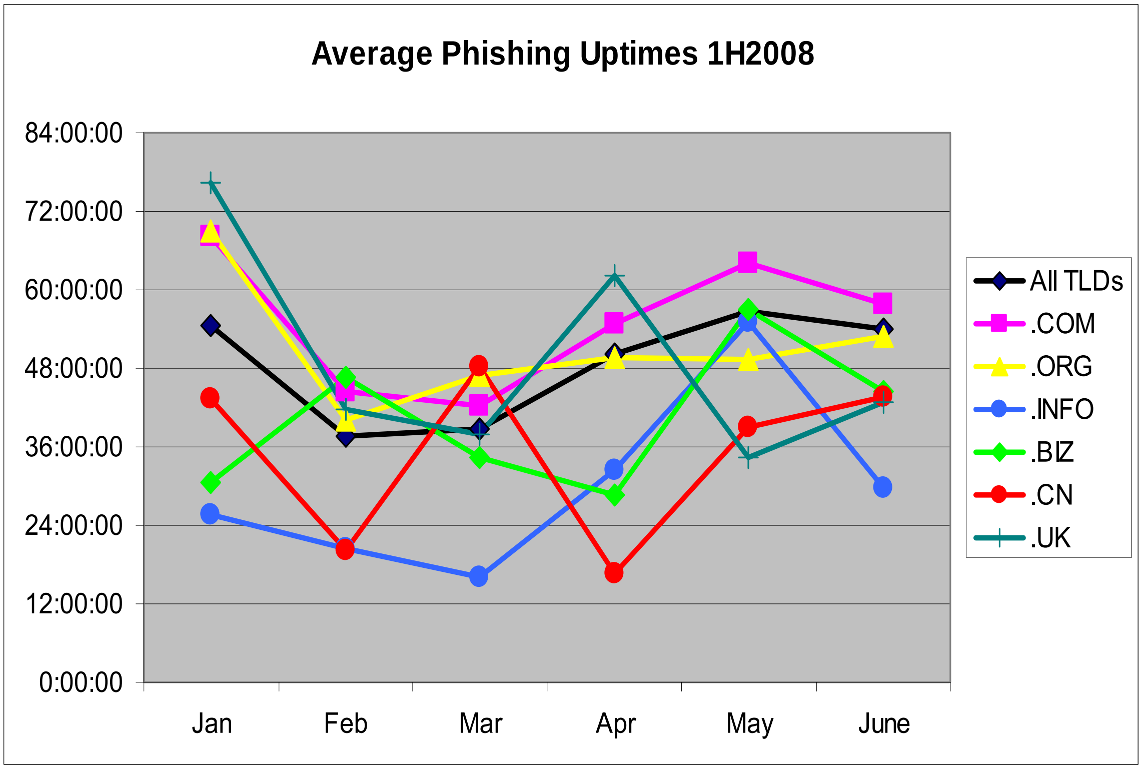

.COM can be used as a reference point; .COM had a slightly longer average uptime than other TLDs, but tracked closely with them. Two gTLDs had notably better performance than the others: .INFO and .BIZ.

Both the .INFO and .BIZ registries are known in the anti-abuse community to have proactive stances for dealing with phishing abuse within their namespaces. Afilias, the .INFO registry operator, had normal uptimes in 2007 and began its anti-phishing program on January 1, 2008. The program cut average uptimes in .INFO by about half, and significantly reduced median uptimes as well.

Another registry that has recently become more proactive is .CN (China, operated by CNNIC). Earlier this year CNNIC announced that it would begin coordinating Chinese registrars in anti-phishing efforts. While .CN average phishing uptimes varied throughout the first half of 2008, they were often significantly better than the Internet average and

.COM. Anecdotally, several Chinese registrars' abuse teams improved their availability and responses to fraudulent registrations utilizing their services as well.

The .INFO, .BIZ, and .CN results seem to show a clear correlation between lower phishing site uptimes and proactive efforts by registry operators and the registrars they work with. In an environment where anti-spam and other security vendors are creating and tuning systems to automatically protect customers from abuse, TLD has become one of several metrics upon which to base the "reputation" of a domain name or URL. So for those service providers who are impacting their abuse statistics, there is a potential pay-off for having their TLDs treated favorably by such systems. This is especially important going into 2009, as ICANN is opening a new round of TLD applications, which will add scores if not hundreds of new TLDs to the Internet. As applicants prepare their business plans and proposals for running new registries, there is compelling evidence that provisions for e-crime response and prevention will have a positive impact.

An even clearer picture of how abuse affects particular TLDs can be seen in the phishing perpetrated by the Rock gang in the spring of 2008. The Rock and several other criminals took advantage of weak anti-abuse policies and procedures in the .HK domain name during 2007. The team at HKDNR worked diligently to change its threat exposure, since the abuse was taking a toll on registry operations and impugning the reputation of the entire .HK domain space. New measures went into place in March 2008, and the number of phishing domains in .HK quickly went from more than 1,000 per month to virtually nothing. Soon thereafter, the Rock started registering and using domain names in the .UK (United Kingdom) and .ES (Spain) TLDs instead. The Rock registered its .UK and .ES names through just a few registrars, and quickly ramped up to register several hundred domain names per month.

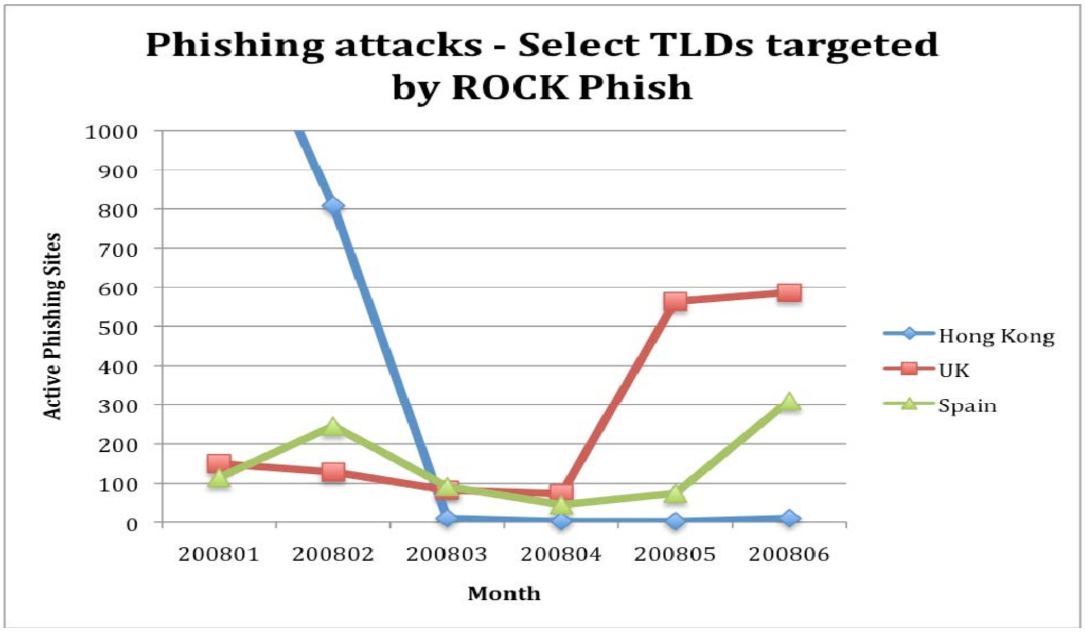

The registrars in question eventually responded, at the behest of the registries and phishing targets.

One of the questions we hoped to answer with this study was whether there was a significant correlation between the number of attacks within a TLD and the length of time those attacks stayed active. We wondered if the number of attacks in a TLD is likely to increase the longer it takes to mitigate sites on domains within that TLD. We also wondered if providers (registrars, brand owners, and registries) who gain experience with lots of phishing attacks would eventually produce shorter take-down times, as the providers involved improve their processes in the face of mounting attacks. We found that there was a slight downward trend in average uptimes and a marked decline in median times the more phish appeared within a particular TLD. This may be due to the work of responders such as brand owners, whose attention is drawn to concentrated attacks.

There was much more variance in outcomes as the number of phish within a TLD got quite small – but with a distinct trend towards longer uptimes. So as plotted on a scatter chart, there is a steep curve downward that flattens to a very long tail:

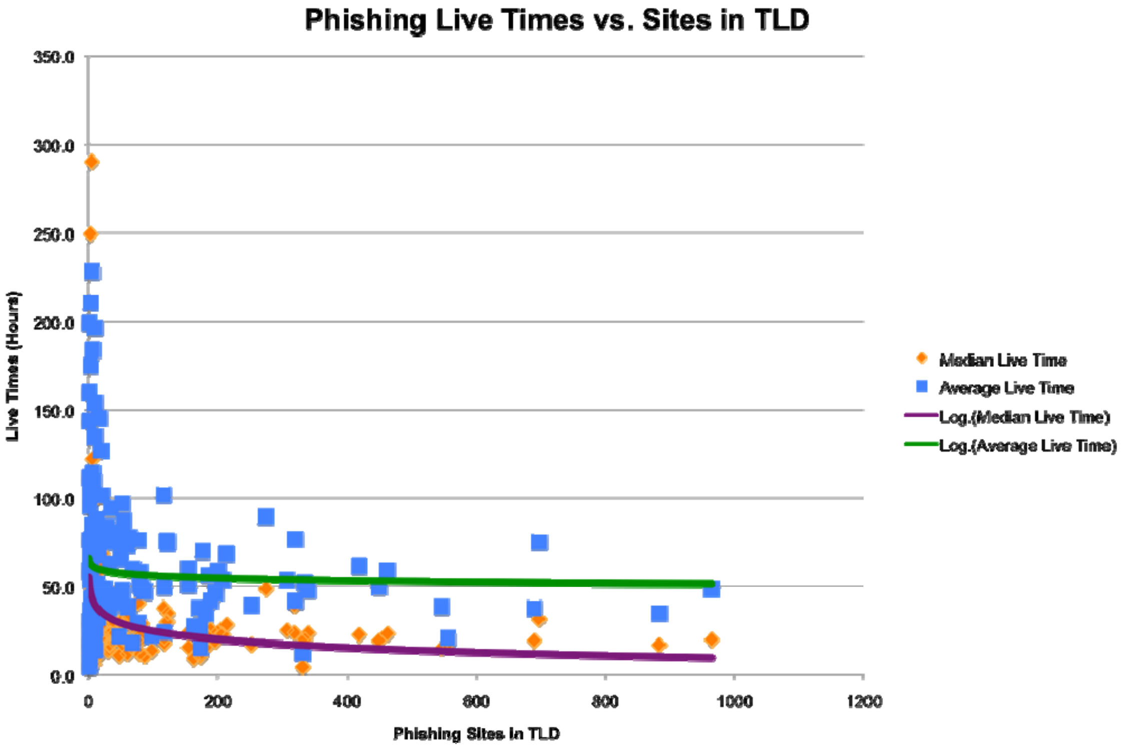

We hope that more data and measurement refinements will help clarify these trends.

Use of Subdomains for Phishing

Use of subdomain registration services is an increasingly worrisome trend, and accounts for the majority of phish in some large TLDs.

We define “subdomain registration services” as providers that give customers subdomain "hosting accounts" beneath a domain name the provider owns. These services offer users the ability to define a "name" in their own DNS space for a variety of purposes. Thus a customer will obtain a hostname to use for his/her own Web site and/or e-mail of the form:

<customer_term>.<service_provider_sld>.TLD

Subdomain registration services include Web hosting companies that provide free subdomain space under their domains, dynamic IP allocation services that supplement their offerings with customizable subdomains, and companies that provide "affinity" subdomains (such as "myfavoriteteam.fan.org"). Some offer DNS services that allow users to redirect their domain names anywhere at any time.

Such services are a popular way for phishers to mount attacks. In our survey we positively identified 4,512 subdomain sites/accounts used for phishing, beneath 274 unique second-level domains. There are likely more within the data set, as it is often difficult to separate them out from other kinds of domains that have hacked hosts or were registered independently by phishers and set up with special subdomains. Even with that caveat, if we had counted these unique subdomains as "regular" domain names, then these types of domains would represent nearly 10% of all domains involved in phishing – a significant percentage.

Top 20 Subdomain Services Used for Phishing 1H2008

	Domain	Phishing Sites	Domain Administrator
1	pochta.ru	379	Pochta.ru
2	land.ru	316	Pochta.ru
3	ns8-wistee.fr	262	wistee.fr
4	9k.com	256	9k.com
5	altervista.org	255	altervista.org
6	smtp.ru	251	Pochta.ru
7	free.fr	250	free.fr
8	nm.ru	171	Pochta.ru
9	t35.com	142	t35.com
10	jexiste.fr	95	jexiste.fr
11	110mb.com	90	110mb.com
12	front.ru	82	Pochta.ru
13	krovatka.su	71	Pochta.ru
14	notlong.com	63	notlong.com
15	freeweb7.com	62	freeweb7.com
16	freehostia.com	60	freehostia.com
17	us.com	55	CentralNIC
18	de.com	45	CentralNIC
19	ifrance.com	44	ifrance.com
20	host.sk	40	host.sk

The Russian freemail provider Pochta.ru owns at least 16 domains that were used to host phishing in 1H2008, and those domains were used to mount at least 1,446 phishing attacks during that period. The subdomains on the single domain pochta.ru were used to target one brand owner beginning on January 17 and running through June, involving at least

379 separate phishing attacks. This underscores how a single provider can host a significant phishing problem.

While much credit can be given to some subdomain providers for quickly mitigating phish on their services, the fact that phishers keep using these services shows that much more anti-abuse work needs to be done. Law enforcement and private investigators observe conversations in the chat rooms and message boards used by the criminal underground, and they emphasize that cybercriminals have very rational, profit-driven approaches. Criminals continue to abuse services where they have success. So the fact that a particular service is used over and over, despite a good post-phish mitigation record, is an excellent indicator that the phishers are making money and will continue to abuse the service until more successful countermeasures are deployed.

The extensive use of subdomain services is eye-opening and poses several challenges. These services are often free, and most are most often offered by small companies. Thus there are few checks and balances on who runs such services or how they screen their customers. These services are typically unmanned or lightly supported, meaning the only point of contact for the domain may be unavailable for days. These conditions are ripe for abuse, both at the consumer level and at the reseller level, as any criminal can set up his own similar service. Depending on the available features of the service, a criminal can obtain as much control over a unique DNS entry as he can through a domain name registrar, making these types of subdomains very convenient for running fast-flux, name-spoofing, and other common domain name tricks used by phishers. There is rarely any published WHOIS information for these subdomains, making it nearly impossible to determine if there is a fraudulent registration, or if someone's legitimate (but hacked) site is being used to host a phish. Instead, responders are completely reliant upon the subdomain service provider to handle all mitigation requests. The fact that there could be thousands of functional, legitimate subdomain sites beneath the main domain means that suspension of the main domain is usually not a viable option. This is an area that invites further research and policy consideration.

Conclusions

This updated study shows that phishers are constantly adapting as they find new opportunities and react to anti-phishing efforts. This study has documented some of their recent strategies and tactics, including their continued abuse of subdomain services, evasion and spoofing techniques, and their systematic exploitation of vulnerable registrars, registries, and subdomain resellers.

Gathering statistics on domain registrations and site uptimes allowed us to show correlations between the efforts of several large gTLD and ccTLD operators and the amount of time phishing sites remained live within their TLDs. The results show that such efforts can lead to significant reduction in the amount of time phishing sites stay live, thus greatly reducing exposure to potential victims of these attacks. We have shown how a single large-scale phishing operation moved from one ccTLD registry to others when forced to.

As we saw in 2007, phishers are engaged in the large-scale use of subdomain services to host and manage their phishing sites. Such attacks can even account for the majority of attacks in certain large TLDs, even to the point of seriously affecting the TLD's overall score in our "Attacks per 10,000 domains" metric.

We see some evidence that the "broken window" theory applies to online service providers. Sociologists created the "broken windows" theory to explain why some neighborhoods thrive, while others decay. The theory posits that ignoring the little problems—graffiti, litter, shattered glass—creates a sense of decline that attracts bad elements and leads the law-abiding to stay away. On the Internet, we wonder if inattentive subdomain providers, registrars, and resellers attract bad actors into a domain space. We have seen some anecdotal evidence of this and it bears watching.

As we noted in our previous study, registrars and registry operators have no control over the security of the Web sites hosted on the domains they sponsor, and have more limited options when vulnerable sites are compromised for phishing. But registries and registrars are in an excellent position to address malicious domain name registrations, which are a major part of the current phishing problem. Registry operators can disseminate information to their registrars, and registrars and hosting providers can mitigate malicious domain name registrations quickly, thereby reducing all phishing up-times and reducing the options available to phishers. We have seen that in action during the period of this study, and what a difference proactive registries and/or registrars can make. We hope this study will spur further research on these and related topics and help the community create improved anti-phishing measures.

Appendix A: Phishing Scores and Up-Times

TLD	TLD Location	# Unique Phishing attacks 1H2008	Unique Domain Names used for phishing 1H2008	Domains in registry in May 2008	Score: Phish per 10,000 domains 1H2008	Score: Attacks per 10,000 domains 1H2008	Average Uptime 1H2008 (HH:MM)	Unique Domain Names used for phishing 2H2007	Domains in registry in Nov 2007	Score: Phish per 10,000 domains 2H2007
ac	Ascension Island	2	1				11:18	2		
ae	United Arab Emirates	9	6				10:11	3		
aero	sponsored TLD	2	2	5,803	**3.4**	**3.4**	12:56	1	5,430	**1.8**
ag	Antigua and Barbuda	2	1	15,005	**0.7**	**1.3**	14:38	1	13,507	**0.7**
al	Albania	2	2	300	**66.7**	**66.7**	5:15	1	250	**40.0**
am	Armenia	9	4	9,686	**4.1**	**9.3**	6:36	11	8,570	**12.8**
ar	Argentina	319	104	1,642,210	**0.6**	**1.9**	11:13	120	1,451,727	**0.8**
as	American Samoa	11	5				10:50	2		
asia	sponsored TLD	2	2	177,707	**0.1**	**0.1**	7:16	0		
at	Austria	176	110	776,150	**1.4**	**2.3**	10:37	158	722,193	**2.2**
au	Australia	307	206	1,098,907	**1.9**	**2.8**	10:19	175	985,458	**1.8**
az	Azerbaijan	3	2	7,100	**2.8**	**4.2**	14:46	1		
ba	Bonnia and Herzegovina	7	7	7,511	**9.3**	**9.3**	6:10	7	6,606	**10.6**
bd	Bangladesh	4	2				9:06			
be	Belgium	690	62	791,737	**0.8**	**8.7**	10:37	163	730,405	**2.2**
bf	Burkina Faso	1	1				23:40	0		
bg	Bulgaria	10	9	8,328	**10.8**	**12.0**	10:27	10	7,500	**13.3**
biz	generic TLD	546	353	2,035,357	**1.7**	**2.7**	9:43	242	1,944,453	**1.2**
bo	Bolivia	6	5	3,900	**12.8**	**15.4**	12:22	3	3,705	**8.1**
br	Brazil	339	204	1,400,423	**1.5**	**2.4**	10:32	317	1,262,967	**2.5**
bs	Bahamas	1	1	1,800	**5.6**	**5.6**	18:49	1		
bt	Bhutan	2	1				12:31	0		
by	Belarus	12	8				13:46	8		
bz	Belize	52	43	43,216	**10.0**	**12.0**	9:03	33	42,360	**7.8**

TLD	TLD Location	# Unique Phishing attacks 1H2008	Unique Domain Names used for phishing 1H2008	Domains in registry in May 2008	Score: Phish per 10,000 domains 1H2008	Score: Attacks per 10,000 domains 1H2008	Average Uptime 1H2008 (HH:MM)	Unique Domain Names used for phishing 2H2007	Domains in registry in Nov 2007	Score: Phish per 10,000 domains 2H2007
ca	Canada	200	146	1,025,000	**1.4**	**2.0**	10:10	157	932,463	**1.7**
cat	sponsored TLD	12	4	29,739	**1.3**	**4.0**	7:13	7	25,885	**2.7**
cc	Cocos (Keeling) Islands	189	91				10:43	98		
cd	Congo, Democratic Republic	2	2				18:11	2		
ch	Switzerland	196	106	1,169,074	**0.9**	**1.7**	10:46	357	1,045,661	**3.4**
cl	Chile	274	128	212,153	**6.0**	**12.9**	10:57	128	195,513	**6.5**
cn	China	1,722	853	11,821,635	**0.7**	**1.5**	10:29	1,540	8,459,174	**1.8**
co	Colombia	68	32	22,303	**14.3**	**30.5**	11:58	34	20,524	**16.6**
com	generic TLD	17,170	12,275	76,625,770	**1.6**	**2.2**	10:16	13,485	70,698,420	**1.9**
coop	sponsored TLD	5	5	5,850	**8.5**	**8.5**	4:45	3		
cr	Costa Rica	6	2	10,830	**1.8**	**5.5**	14:03	2	6,905	**2.9**
cx	Christmas Island	60	45	4,539	**99.1**	**132.2**	9:47	6	4,387	**13.7**
cy	Cyprus	5	3	7,875	**3.8**	**6.3**	11:53	1	8,229	**1.2**
cz	Czech Republic	116	74	432,246	**1.7**	**2.7**	11:37	105	354,592	**3.0**
de	Germany	964	711	12,072,501	**0.6**	**0.8**	11:02	903	11,524,091	**0.8**
dj	Djibouti	1	1				15:55	0		
dk	Denmark	117	87	913,000	**1.0**	**1.3**	10:56	128	862,000	**1.5**
dm	Dominica		0					1	19,469	**0.5**
ec	Ecuador	20	14	16,250	**8.6**	**12.3**	12:36	10	14,941	**6.7**
edu	U.S. higher education	34	19	7,000	**27.1**	**48.6**	11:25	34	6,997	**48.6**
ee	Estonia	16	11	58,241	**1.9**	**2.7**	11:32	18	51,831	**3.5**
eg	Egypt	2	2	4,500	**4.4**	**4.4**	15:28	1	4,490	**2.2**
es	Spain	883	333	970,580	**3.4**	**9.1**	9:58	181	770,984	**2.3**
eu	European Union	213	134	2,737,047	**0.5**	**0.8**	11:23	117	2,671,846	**0.4**
fi	Finland	32	22	180,926	**1.2**	**1.8**	10:23	18	165,000	**1.1**
fk	Falkland Islands	1	1				7:33	0		
fm	Micronesia, Fed. States	3	2				7:50	2		
fr	France	1,236	107	1,128,776	**0.9**	**10.9**	11:05	168	969,864	**1.7**
ge	Georgia	6	4	61,292	**0.7**	**1.0**	7:20	3	7,199	**4.2**
gg	Guernsey	6	6				7:00	1	1	

TLD	TLD Location	# Unique Phishing attacks 1H2008	Unique Domain Names used for phishing 1H2008	Domains in registry in May 2008	Score: Phish per 10,000 domains 1H2008	Score: Attacks per 10,000 domains 1H2008	Average Uptime 1H2008 (HH:MM)	Unique Domain Names used for phishing 2H2007	Domains in registry in Nov 2007	Score: Phish per 10,000 domains 2H2007
gh	Ghana	3	2				6:39	3		
gov	U.S. government	2	2				7:52	1	1	
gp	Guadeloupe		0					1		
gr	Greece	82	60	203,000	**3.0**	**4.0**	12:50	44	202,000	**2.2**
gs	South Georgia and South Sandwich Islands	1	1	8,400	**1.2**	**1.2**	6:35	4	8,300	**4.8**
gt	Guatemala	9	5	7,070	**7.1**	**12.7**	10:16	4	6,262	**6.4**
hk	Hong Kong	2,278	516	160,336	**32.2**	**142.1**	10:59	1,024	148,757	**68.8**
hm	Heard and McDonald Is.		0					2		
hn	Honduras	8	3	4,314	**7.0**	**18.5**	14:21	16	3,820	**41.9**
hr	Croatia	20	12	58,779	**2.0**	**3.4**	11:26	12	51,432	**2.3**
ht	Haiti	1	1	1,125	**8.9**	**8.9**	17:12	0		
hu	Hungary	120	85	372,700	**2.3**	**3.2**	11:27	82	350,091	**2.3**
id	Indonesia	49	30				12:15	35		
ie	Ireland	15	12	103,168	**1.2**	**1.5**	14:13	15	90,710	**1.7**
il	Israel	32	20	126,303	**1.6**	**2.5**	10:47	25	112,500	**2.2**
in	India	252	150	454,330	**3.3**	**5.5**	10:41	85	331,495	**2.6**
info	generic TLD	1,430	684	5,042,032	**1.4**	**2.8**	9:24	499	4,956,218	**1.0**
io	British Indian Ocean Terr.	29	14				7:55	5		
ir	Iran	32	21	89,890	**2.3**	**3.6**	10:42	12	72,947	**1.6**
is	Iceland	14	5	22,000	**2.3**	**6.4**	10:14	8	20,000	**4.0**
it	Italy	462	223	1,510,009	**1.5**	**3.1**	11:11	218	1,467,221	**1.5**
je	Jersey	3	1				5:41	0		
jp	Japan	207	152	1,020,763	**1.5**	**2.0**	10:58	177	972,584	**1.8**
ke	Kenya	3	2	9,193	**2.2**	**3.3**	10:22	3	8,011	**3.7**
kg	Kyrgyzstan	76	38	2,500	**152.0**	**304.0**	11:33	21		
kh	Cambodia	3	2				10:52	1		
kr	Korea	697	345	945,000	**3.7**	**7.4**	11:04	240	932,841	**2.6**
ky	Cayman Islands	1	1	5,600	**1.8**	**1.8**	5:10	0		

TLD	TLD Location	# Unique Phishing attacks 1H2008	Unique Domain Names used for phishing 1H2008	Domains in registry in May 2008	Score: Phish per 10,000 domains 1H2008	Score: Attacks per 10,000 domains 1H2008	Average Uptime 1H2008 (HH:MM)	Unique Domain Names used for phishing 2H2007	Domains in registry in Nov 2007	Score: Phish per 10,000 domains 2H2007
kz	Kazakhstan	7	6				5:50	7		
la	Lao People's Democratic Rep.	79	28				11:23	15		
li	Liechtenstein	97	26	59,546	**4.4**	**16.3**	11:12	218	50,100	**43.5**
lk	Sri Lanka	5	2				12:20	4		
lt	Lithuania	24	18	56,512	**3.2**	**4.2**	13:28	37	64,554	**5.7**
lu	Luxembourg	8	6	40,500	**1.5**	**2.0**	12:35	6	34,000	**1.8**
lv	Latvia	14	8	30,000	**2.7**	**4.7**	11:40	19	28,900	**6.6**
ly	Libya	39	39	3,181	**122.6**	**122.6**	13:42	0	3,100	**0.0**
ma	Morocco	5	4	24,500	**1.6**	**2.0**	9:27	2	25,873	**0.8**
md	Moldova	6	6				7:42	9	2,200	**40.9**
mk	Macedonia	1	1				6:22	5		
mn	Mongolia	68	20	6,378	**31.4**	**106.6**	10:12	89	5,087	**175.0**
mo	Macao	5	3	2,200	**13.6**	**22.7**	10:43	3		
mobi	sponsored TLD	173	85	918,634	**0.9**	**1.9**	9:25	45	761,549	**0.6**
ms	Montserrat	47	38				10:28	4		
mx	Mexico	122	83	255,406	**3.2**	**4.8**	11:36	99	230,177	**4.3**
my	Malaysia	19	11	117,718	**0.9**	**1.6**	14:15	19	98,000	**1.9**
name	sponsored TLD	331	126	289,343	**4.4**	**11.4**	5:25	38	265,638	**1.4**
net	generic TLD	4,159	2,305	11,623,856	**2.0**	**3.6**	10:06	2,106	10,581,849	**2.0**
nf	Norfolk Island	1	1				10:19	0		
ng	Nigeria	1	1				21:12	1		
ni	Nicaragua	2	2	4,600	**4.3**	**4.3**	17:33	1	4,254	**2.4**
nl	Netherlands	449	305	2,919,646	**1.0**	**1.5**	10:35	413	2,661,308	**1.6**
no	Norway	76	60	387,238	**1.5**	**2.0**	10:31	51	357,722	**1.4**
np	Nepal	3	3	11,865	**2.5**	**2.5**	10:11	10	11,016	**9.1**
nr	Nauru	2	2				16:13	1		
nu	Niue	181	93				10:20	28		
nz	New Zealand	32	26	332,794	**0.8**	**1.0**	9:04	63	311,198	**2.0**
org	generic TLD	2,384	1,425	6,905,011	**2.1**	**3.5**	10:28	1,488	6,412,064	**2.3**
pa	Panama	6	2	4,600	**4.3**	**13.0**	11:46	1	4,488	**2.2**
pe	Peru	49	24	25,714	**9.3**	**19.1**	11:50	17	17,859	**9.5**

TLD	TLD Location	# Unique Phishing attacks 1H2008	Unique Domain Names used for phishing 1H2008	Domains in registry in May 2008	Score: Phish per 10,000 domains 1H2008	Score: Attacks per 10,000 domains 1H2008	Average Uptime 1H2008 (HH:MM)	Unique Domain Names used for phishing 2H2007	Domains in registry in Nov 2007	Score: Phish per 10,000 domains 2H2007
ph	Philippines	170	63				9:02	164		
pk	Pakistan	7	6				12:13	9		
pl	Poland	417	300	960,000	**3.1**	**4.3**	11:16	296	753,520	**3.9**
ps	Palestinian Territory	9	4				9:21	5		
pt	Portugal	37	22	215,000	**1.0**	**1.7**	9:24	42	184,596	**2.3**
py	Paraguay		0	7,700				5	6,501	**7.7**
ro	Romania	184	142	284,700	**5.0**	**6.5**	11:04	155	242,484	**6.4**
ru	Russia	1,907	362	1,427,928	**2.5**	**13.4**	11:44	360	1,104,572	**3.3**
sa	Saudi Arabia	10	6	13,160	**4.6**	**7.6**	6:33	6	12,478	**4.8**
se	Sweden	79	50	737,000	**0.7**	**1.1**	10:18	64	685,000	**0.9**
sg	Singapore	15	11	94,895	**1.2**	**1.6**	10:03	13	87,086	**1.5**
sh	Saint Helena	2	2				11:56	2		
si	Slovenia	7	6	56,000	**1.1**	**1.3**	10:55	9	50,312	**1.8**
sk	Slovakia	63	50	159,758	**3.1**	**3.9**	10:59	55	150,601	**3.7**
sl	Sierra Leone	1	1				4:21	0		
st	Sao Tome and Principe	556	122				8:46	40		
su	Soviet Union	154	14	60,543	**2.3**	**25.4**	12:11	12	19,431	**6.2**
sv	El Salvador	0	0	3,900	**0.0**	**0.0**		3	4,184	**7.2**
tc	Turks and Caicos	60	44	10,300	**42.7**	**58.3**	10:30	5	9,000	**5.6**
th	Thailand	154	84	35,757	**23.5**	**43.1**	11:35	99	33,000	**30.0**
tj	Tajikistan	12	6				13:50	2		
tk	Tokelau	162	104	1,310,000	**0.8**	**1.2**	8:46	29	1,869,000	**0.2**
tl	Timor-Leste	9	5				12:18	2		
tm	Turkmenistan	5	4				2:36	0		
tn	Tunisia	2	1				22:56	3		
to	Tonga	54	32				9:37	14		
tp	Portuguese Timor	1	1				11:07	1		
tr	Turkey	32	26	156,390	**1.7**	**2.0**	10:09	37	142,646	**2.6**
travel	sponsored TLD	1	1	201,294	**0.0**	**0.0**	6:52	1	28,665	**0.3**
tt	Trinidad and Tobago	1	1	2,100	**4.8**	**4.8**	15:26	0		
tv	Tuvalu	62	32				9:34	81		

ua	Ukraine	123	66	359,377	**1.8**	**3.4**	11:47	51	311,822	**1.6**
uk	United Kingdom	1,584	922	6,830,000	**1.3**	**2.3**	10:24	529	6,445,465	**0.8**
us	United States	334	256	1,419,000	**1.8**	**2.4**	9:55	271	1,362,805	**2.0**
uy	Uruguay	9	5	14,795	**3.4**	**6.1**	12:10	7	13,936	**5.0**
uz	Uzbekistan	3	3				10:59	1		
ve	Venezuela	86	71	75,000	**9.5**	**11.5**	8:00	22	53,704	**4.1**
vg	British Virgin Islands	117	30	8,138	**36.9**	**143.8**	10:34	1	7,405	**1.4**
vi	Virgin Islands	1	1				4:10	0		
vn	Vietnam	28	17	69,473	**2.4**	**4.0**	9:54	9	54,739	**1.6**
vu	Vanuatu		0					6		
ws	Samoa	173	97	588,000	**1.6**	**2.9**	9:03	48	522,221	**0.9**
yu	Yugoslavia	5	5	50,370	**1.0**	**1.0**	5:41	13	46,279	**2.8**
za	South Africa	52	38	394,749	**1.0**	**1.3**	10:45	64	359,518	**1.8**
zm	Zambia	10	6				11:38	0		
zw	Zimbabwe	5	3				7:58	2		

About the Authors

Greg Aaron is Director of Key Account Management and Domain Security at Afilias (www.afilias.info). Afilias operates the .INFO top-level domain (TLD) and provides technical and advising services for thirteen other TLDs, including .ORG, .MOBI, .ASIA, .ME, and .IN (India). Greg oversees .INFO operations and Afilias' security programs, including domain name abuse policy and practices. He is also an expert on domain name intellectual property issues and Internationalized Domain Names (IDNs). He serves on the steering committee of the Anti-Phishing Working Group (APWG), has advised the Government of India regarding domain and related Internet policies, and is an active member of ICANN's Fast-Flux Working Group. He previously worked at Internet companies such as Travelocity, and graduated magna cum laude from the University of Pennsylvania.

Rod Rasmussen is President and CTO of Internet Identity (www.internetidentity.com), and has served as its technical leader since he co-founded the company in 2001. He is widely recognized as a leading expert on the abuse of the domain name system by phishing criminals. Rasmussen is co-chair of the Anti-Phishing Working Group's (APWG) Internet Policy Committee (IPC), and serves as the APWG's Industry Liaison to various groups around the world, including ICANN, the international oversight body for domain names. He serves on ICANN's Fast-Flux Working Group. He is also a member of the Steering Committee for the Authentication and Online Trust Alliance (AOTA), and an active member of the Digital PhishNet, a collaboration between industry and law enforcement. Prior to starting Internet Identity, Rasmussen held product management roles for LanQuest, a network equipment testing company, and networking product manufacturer Global Village. Rasmussen earned an MBA from the Haas School of Business at the University of California, Berkeley and holds two bachelor's degrees, in Economics and Computer Science, from the University of Rochester.

#