APWG
Internet Policy Committee

Global Phishing Survey:
Domain Name Use and
Trends in 2007

Global Phishing Survey:

Domain Name Use and Trends in 2007

Greg Aaron

Afilias

<gaaron at afilias.info>

Rod Rasmussen

Internet Identity

<rod.rasmussen at internetidentity.com>

May 26, 2008

Table of Contents

Summary

In order to combat phishing effectively, it is important to understand how phishers use domain names and to what purposes. Domain name usage is an important measure of the scope of the global phishing problem, and understanding why and how phishers register domain names for their own use can lead to improvements in anti-abuse measures. Analysis of URL construction provides clues about how phishers mislead Internet users, and reveals how phishers are using certain online service providers.

This study describes our analysis of a comprehensive database of the phishing that took place in 2007. Specifically, the data includes all the phishing attacks detected between January 1, 2007 and December 31, 2007 that were collected by the APWG, supplemented with additional reports from several phishing feeds and private sources. The APWG phishing attack repository is the Internet's most comprehensive archive of e-mail fraud and phishing activity. The data set includes the URLs of the attacks, and their targets. Our study is designed to complement rather than duplicate the APWG's monthly 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. [1]

Our data reveals many interesting ways that phishers use domain names in their attacks. Some are common knowledge within the anti-phishing community, but others are surprising, and we hope that bringing these tactics to light will lead to improved anti-phishing measures.

Our major findings are summarized below:

1. We have created a metric to measure the prevalence of phishing in different top-level domains. This metric provides one way to compare top-level domains (TLDs) to each other, and shows that the pervasiveness of phishing varies greatly across TLDs. The metric is a useful tool for identifying TLDs that have been exploited by phishers who register domain names.

[1] These reports are available at: http://www.apwg.org/phishReportsArchive.html

2. The domain name used for phishing (and therefore the TLD) rarely matters to phishers. The vast majority of domain names used for phishing do not contain a brand name, and many are meaningless in and of themselves. Instead, phishers increasingly embedded brand names in subdirectories or subdomains.
3. Phishers choose certain TLDs to register domain names in, and change their preferences over time. The phishers choose based on the availability and policies of the different TLD registries, the services and TLDs offered by its registrars, and the anti-abuse practices of the registrars and the registry operator. Domain name prices seem to be a secondary factor at best.
4. Phishers are engaged in the large-scale use of subdomain hosting services for phishing sites, and there is a continued reliance on hacked or compromised Web sites. There appears to be a waning use of well-known shared hosting services, where fake accounts can be set up as subdirectories on trusted hosting domains.
5. Domain name registrars and registries are in a good position to monitor and mitigate domain name registrations made by phishers.

Basic Statistics

Millions of phishing URLs were reported in 2007, but the number of phishing attacks and domain names used to host them is much smaller. This is due to several factors:

1. Some phishing involves customized attacks that track targeted individuals or groups through a numbering system of some sort in the URL. What is basically one phishing attack is therefore represented as many URLs, sometimes one for each spam e-mail sent by the phisher. Transmitting unique URLs confuses spam filters looking for repeated links, fools collators into recording duplicate entries, and misleads blacklist users who search for exact matches.
2. Phishers often use one domain name to host simultaneous attacks against multiple brands. For example, the Rock Phish gang often hosts five or more phishes on one domain name.
3. A phishing site may have multiple pages, each of which may be reported.

These factors complicate any discussion of the number of phish per domain name.

The data set yielded the following basic statistics:

- The attacks utilized 51,989 unique domain names.[2]
- In addition, 15,715 phish were found on IP addresses rather than on domain names. (For example: http://91.121.81.84/do.php?cmd=SignIn.) 11,553 of the IP addresses used were unique.
- Phishing took place on domain names in 182 TLDs. This is two-thirds of the 273 TLDs in existence.
- Only 12 of the 51,989 domain names were Internationalized Domain Names (IDNs). Of the 12, nine were in .HK, and the other three were in .DE.

Each domain name's registrar of record was often 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 by a new registrant and/or registrar. Reconstructing point-in-time registrar sponsorship of a domain name often requires registry-level data. This historical data is usually closely held, and it is not practical to obtain it from all the registries involved. Registrar-specific statistics and trends are certainly of interest, and are an opportunity for future studies.

Compromised Domains, Malicious Registrations, and URL Construction

Phishing often takes place on compromised computers, where the phishers place their phishing pages unbeknownst to the site operators. This method gains the phishers free hosting, and complicates take-down efforts because suspending a domain name or hosting account also disabled the resolution of the legitimate user's site. Phishing on a compromised Web site typically takes place on a subdomain or in a subdirectory, where the phish is not easily noticed by the site's operator or visitors. A common trick is to use a leading period (".") in the name of the directory the phishing site is stored in, making it difficult to find with standard directory listing commands.

[2] "Domain names" are defined as second-level domain names, plus third-level domains 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.).

Phishers often embed a relevant brand name in the URL in order to fool those lured to the phish.[3] The first example below is of a phish and brand name embedded in a subdirectory of a .CA name, while the second shows a brand name placed on a subdomain of a .NET name:

- http://www.domainname.ca/~test/bankname/login/signon.htm
- http://www.bankname.com.447956.33njm34webnyq2.net/cmd-confirm/login.php

This obfuscation sometimes makes it difficult to determine whether a given domain name was registered by a phisher. Also, some domain names used for phishing are obtained on the secondary market, and registrant data in WHOIS is often faked or obscured by proxy services.

Given these caveats, we identified 10,773 of the 51,989 domain names in our data set as "malicious" registrations made by phishers, specifically registered to host phishing sites. This 20% figure is conservative, and we believe that the percentage of malicious registrations is actually much higher. Maliciously registered domains were identified as such if they were reported for phishing within a very short time of being registered (this is an indicator that their sites were not compromised), or were registered in batches or in patterns that indicated common ownership or intent.

The data reveals that phishers prefer to obscure the base domain name even when the phisher has registered a domain name for his own use. Of the 10,773 maliciously registered domains, 10,515 had their phishes placed on subdomains or in subdirectories. This served at least two purposes:

1) It allowed the phisher to use a domain name that did not contain a brand name or variant thereof. Instead, the brand name was contained somewhere else in the URL.
2) In many instances, use of subdomains or subdirectories allowed the phishers to embed several different phishes on one domain name, targeting different brands. This method is routinely employed by the Rock Phish gang.

The other 258 domains had phishes that appeared on the "base domain" level or "home" page of a second-level domain (http://www.baddomain.tld or

[3] This trend accelerated in 2007. According to the APWG's monthly reports, the percentage of phish containing some form of target name in the URL rose from 25% early in the year to a high of 42.1% in December 2007.

http://baddomain.tld). 201 of those domains were in the .COM TLD. Approximately half of those 258 domains contained a brand name or variant thereof, and were designed to fool visitors by looking legitimate. Typical formats included:

- http://admin-bankname.tld
- http:// bankname-verification.tld
- http://www.brandonlineupdate.tld

The other half did not contain a brand name or other enticement. Many were random strings, such as http://rjt27.com, that offered nothing to confuse a potential victim.

We observed no meaningful correspondence between the home country of the brands being targeted and the TLD of the domains used to attack them. .COM was the TLD used most often to target the most-phished brands, and .NET the second-most-used. In one example, 359 .HK domains and 200 .CN domains were used to target a prominent German bank—but only two .DE names were used to phish that target.

Our conclusions are that:

1. There was an increase in the practice of embedding a brand name or other misleading string somewhere in the URL in order to fool victims.
2. Phishers continue to use compromised domains, and the practice is apparently effective.
3. Phishers also register a significant number of domains for their own use. See "The Rock Phish Factor" below for related commentary.
4. The domain name itself usually does not matter to phishers. Therefore a domain name in any TLD will do. Internationalized domain names (IDNs) represented a miniscule percentage of the domains used.
5. Brand name owners should continue to make defensive domain name registrations, and should continue to use detection methods that find infringing domain names by scanning zone files for pattern matches. However, the data indicates that phishers are probably aware of that countermeasure and avoid domain names that draw attention to themselves. Brand owners should also employ detection methods that collect and analyze entire phishing URLs.

Use of Subdomain Registration Services for Phishing

We define "subdomain registration services" as companies that provide customers with free or paid-for subdomain "hosting accounts" beneath the company's own domain name(s). These services provide users with 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 those that provide "affinity" subdomains (such as "myfavoriteteam.fan.org"), Web hosting companies that provide free subdomain space under their domains, and dynamic IP allocation services that supplement their offerings with customizable subdomains. Some offer DNS services that allow users to redirect their domain names anywhere at any time.

In our survey we positively identified 11,443 subdomain sites/accounts used for phishing, beneath 448 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 at least 18% of all domains involved in phishing – a significant percentage.

Examples of subdomain accounts used for phishing from our survey data include:

- account-slgnln-elbay-fr.pochta.ru. (Pochta.ru is a popular free e-mail service that offers unlimited mailboxes and free hosting.)
- labsupport.no-ip.org. (The domain no-ip.org redirects to No-IP.com, a company that provides managed DNS, dynamic DNS, domain registration, e-mail, and other domain-related services.)
- A free online tool that makes it easy for anyone to create and publish Web pages in just minutes. This service hosted multiple phishes that targeted social networking sites, an auction provider, and other brands in 2007.

We observed waning use of hosting services that offer subdirectory accounts rather than subdomains.[4] The major providers of such services may be more vigilant than they once were, and subdomains may be more versatile for phishing.

The extensive use of subdomain services is eye-opening and poses several challenges. These services are unaccredited (unlike domain name registrars are), are often free, and most are offered by small companies. Thus there are few checks and balances on

[4] An example of such as phish was: http://geocities.yahoo.com.br/orkuttcomunittaspp/Orkut.Com.htm

who runs such services or how they screen their customers. These conditions are ripe for abuse, both at the consumer level and at the reseller level, as any criminal can set up his own such 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 no 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. In the latter case, the lack of WHOIS makes it much harder to track down the site owner of a hacked Web site during a take-down effort.

Instead, responders are completely reliant upon the subdomain service provider to handle all mitigation requests. These services are typically unmanned or lightly supported, meaning the only point of contact for the domain may be unavailable for days. 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.

Prevalence of Phishing by Top-Level Domain (TLD)

We then sorted the 51,989 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. So to create a different basis for comparison and to place the numbers in context, we developed a metric that measures the *prevalence* of phishing in a TLD.

This metric – "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.[5] This metric is a way of revealing whether a TLD has a higher or lower incidence of phishing relative to others. Based on this metric, some interesting trends reveal themselves.

[5] Score = (phishing domains / domains in TLD) x 10,000

Phishing occurred on domain names in 182 TLDs. Of these, we were able to obtain the domain count statistics for 105 TLD registries.[6] These 105 TLDs contained 97.6% of the phishing domains in our data set (50,774 out of the 51,989), and a total of 150,689,751 domain names overall. Industry estimates put the total number of domain names in existence worldwide at the end of 2007 at approximately 153,000,000.[7]

The complete tables are presented in Appendices A and B, including the scores and the number of phishes in those 182 TLDs.

- The median score was 4.7.
- The average score was 15.3, which was skewed by a few high-scoring TLDs.
- The standard deviation was 35.3, with a confidence level (95.0%) of 6.8.
- .COM, the world's largest and most ubiquitous TLD, had a score of 3.4. .COM contains 45.9% of the phishing domains in our data set, and 46.9% of the domains in the TLDs for which we have domains-in-registry statistics. In the ranking of TLDs by score, there are 46,402,669 domains in the TLDs ranked below .COM, and 33,597,962 in the TLDs ranked above .COM.

We therefore suggest that scores between .COM's 3.4 and the median 4.7 occupy a middle ground, with scores above 4.7 indicating TLDs with increasingly prevalent phishing.

Notes regarding the statistics:

- 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 even one phisher, or the inattention of one registrar. (See "Factors Affecting Phishing Scores: Registrars" and "Effect of and Response to Malicious Registrations" below for related notes.)

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

[6] For the purposes of this study, we used the number of domain names in each registry as of November 2007. Sources: ICANN.org (for gTLD and sTLD monthly registry reports), ccTLD registry operators, Latinoamericann.org.
[7] VeriSign and Zooknic, http://www.verisign.com/static/043379.pdf

Top 20 Phishing TLDs in 2007 by Score

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

Rank	TLD	TLD Location	Domains in registry in November 2007	Domain names used for phishing in 2007	Score: Phish per 10,000 domains
1	.hk	Hong Kong	150,799	1,707	**113.2**
2	.th	Thailand	33,000	171	**51.8**
3	.li	Liechtenstein	50,100	221	**44.1**
4	.ro	Romania	242,484	316	**13.0**
5	.cl	Chile	195,513	222	**11.4**
6	.bz	Belize	42,360	48	**11.3**
7	.tw	Taiwan	341,462	361	**10.6**
8	.lt	Lithuania	64,554	65	**10.1**
9	.ee	Estonia	50,000	47	**9.4**
10	.cz	Czech Republic	347,989	286	**8.2**
11	.mx	Mexico	230,177	189	**8.2**
12	.pl	Poland	753,520	581	**7.7**
13	.sk	Slovakia	150,601	107	**7.1**
14	.ve	Venezuela	53,704	36	**6.7**
15	.yu	Yugoslavia	46,279	30	**6.5**
16	.ru	Russia	1,104,572	684	**6.2**
17	.at	Austria	722,193	415	**5.7**
18	.tr	Turkey	142,646	73	**5.1**
19	.in	India	331,495	168	**5.1**
20	.hu	Hungary	350,000	173	**4.9**

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

Phishing in gTLDs in 2007 by Score

Rank	TLD	Domains in registry November 2007	Domain names used for phishing in 2007	Score: Phishing domains per 10,000
65	.org	6,412,064	2,627	**4.1**
68	.biz	1,944,453	764	**3.9**
70	.net	10,581,849	3,973	**3.8**
79	.com	70,698,420	23,860	**3.4**
88	.info	4,954,266	1,295	**2.6**

The sizeable TLDs with the lowest scores were:

rank	TLD	TLD Location	Domains in registry in November 2007	Domain names used for phishing in 2007	Score: Phishing domains per 10,000
94	.cn	China	8,459,174	1,853	**2.2**
95	.ws	Samoa	522,221	114	**2.2**
96	.name	sponsored TLD	265,638	55	**2.1**
97	.se	Sweden	685,000	127	**1.9**
98	.ar	Argentina	1,451,727	230	**1.6**
99	.de	Germany	11,524,091	1,798	**1.6**
100	.uk	United Kingdom	6,445,465	992	**1.5**
102	.eu	European Union	2,671,846	197	**0.7**
103	.mobi	sponsored TLD	761,549	48	**0.6**

Factors Affecting Phishing Scores

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?

Registrant Base

Are certain TLDs susceptible to phishing because their registrants have worse Web site security and therefore suffer more site compromises? Registrant base does not seem to account for the variance in gTLD scores. The highest-scoring gTLD was .ORG, an open TLD that is generally used by and is associated with noncommercial entities. But close behind was .BIZ, which is marketed for business use. The lowest phishing incidence among gTLDs was in .INFO, which is an open TLD with a mix of registrants.

Regarding ccTLDs, it is possible that Web site security is less adequate than average in certain countries. But the theory is not always supported by the numbers. For example, Austria and Germany are neighbors who are on par technically and economically, but .AT had a score of 5.7 while .DE scored a low 1.6.

.EDU had the fourth-highest score of any TLD -- a 95.8 – and appears to be a special case. Registration of .EDU domains is carefully regulated, and so all .EDU phish were the result of site compromises. A typical university Web site may be vulnerable because it is often a sprawling affairs managed by different schools or departments using various subdomains, content management systems, and hosting platforms.

Price

Phishers have the means to register domains in the TLDs of their choice, regardless of the retail price. They are in the business of stealing financial instruments, and often have a supply of stolen credit card numbers that they can use to illegitimately register domain names.

Price does not seem to account for the variance in gTLD scores. All the gTLD registries offer their domains at similar wholesale prices (around US$6.15) and are generally sold at competitive retail prices across the same registrars. They occasionally offer sales specials and bulk discounts to their registrars.

In March 2007, the .CN registry operator, CNNIC, significantly reduced the annual cost of .CN domain name registrations to one yuan (US$0.13). The low price helped .CN grow explosively, from 1.87 million domains in February 2008 to 9 million in December

2008.[8] However, the decrease in price did not lead to an immediate increase in demand by phishers. Phishing in .CN remained at relatively low levels until August 2007. At that point phishers seem to have "discovered" .CN—or determined it advantageous for their purposes—and began registering .CN domains for their own use. 60% of the year's phishes on .CN names took place in the last four months of the year.

Domain Usage Rates

The number of active domains in each TLD may be a factor. A certain percentage of domains in any TLD either do not resolve, or do not host unique content. Domains without such content are usually not as vulnerable to compromise.

- Many domain names are "parked" at placeholder pages supplied by registrars or hosting providers. The .COM domain is especially home to networks of "pay-per-click" pages meant to monetize domain names that Internet users find via direct navigation or through typographical errors. Parking and pay-per-click pages are supported by hosting and content management systems that either have a decent level of security, or were not hacked by phishers.
- It takes two to three years for a new TLD to build up actual Web site usage, and in such TLDs there are fewer active, resolving domains that phishers can compromise. Names that do resolve in a new TLD are often simply redirected to the owner's pre-existing site on another TLD. These factors may help explain the very low phishing scores of the .EU and .MOBI TLDs, which were young in 2007.[9]

Ease of Registration: Registry Policy and Technology

Many phishers prefer to register domains that offer easy online registration and rapid DNS updates, which they can use to launch attacks within minutes or hours of registration. This helps maximize phishing site up-time, especially if the phisher uses stolen credit card information that may trigger anti-fraud alarms.

Malicious registrations can be curbed if the TLD registry limits the availability of its domains to qualified parties. These policy impediments take the form of residence, citizenship, or other "nexus" requirements imposed by ccTLDs, or the "community" or affinity requirements imposed by some sponsored top-level domains (sTLDs). These barriers are enforced in various ways, sometimes at the time of registration, and some

[8] While the prevalence of phishing in .CN remained low in 2007, there were reports of increased cybersquatting due to the .CN price decrease. See http://www.fairwindspartners.com/perspectives-vol-02-issue-05.html and http://www.news.com/Cybersquatting-escalates-in-Asia/2100-1030_3-6212187.html

[9] The .EU Land Rush took place in April 2006. The .MOBI Land Rush took place in October 2006.

present a delay or waiting period. The .IE (Ireland) registry is an example. It only began allowing domain names for personal use in October 2007, and requires qualified applicants to provide documentary evidence of the applicant's legal name, such as a copy of the applicant's passport or birth certificate.

Registries that use proprietary registration technologies or protocols are often available through a smaller number of registrars, or a set of localized registrars that may not offer a wide range of TLDs. For example, some of the world's largest global registrars do not sell certain large ccTLDs because they feel that those registries' proprietary technologies do not justify the setup and maintenance costs. Large ccTLDs tend to have most of their registrars located in-country, and this is the case for .DE, .UK, .NZ, .AU, .IE, and others.

In other words, domain name registries have choices about how and to whom they will offer their domain names, and those choices involve trade-offs between convenience and risk.

Finally, a registry's anti-abuse process can also make a significant difference. A registrant's main business relationship is with the registrar. Many registries therefore push abuse reports to the registrar for investigation and follow-up. This process takes time, and extends the up-time of phishing sites.

Registrars

Domain name registrars come in all sizes and levels of ability. Their ability and willingness to respond to abuse reports varies widely. Many are small companies, and are only loosely overseen by the bodies (ICANN and the registry operators) that accredit them. It takes only one inattentive or irresponsible registrar to allow a batch of malicious registrations, and thereby create a large problem for anti-abuse responders worldwide.

We and other researchers have observed that the Rock Phish gang perpetrates a large number (perhaps the majority) of malicious domain registrations. The gang's "business model" is to productionize and launch large numbers of attacks on a regular basis. This requires a large number of domain names, so the gang simply registers what it needs.

The Rock Phish gang will often attack previously untargeted registrars who are slow to respond, or have weak credit card authentication. If the registrars realize what has happened and put effective processes in place to suspend the domains quickly, the gang moves on to previously untargeted registrars. Interestingly, the gang minimizes its risk by spreading its registrations across TLDs, sometimes registering the same string in an array of TLDs on the same day, via the same registrar.

In their excellent examination of phishing take-downs[10], Tyler Moore and Richard Clayton made a related observation. Asking why the Rock Phish gang continued to buy and activate new domains even when their earlier ones still worked, they noted: "One reason is that the domains may lose effectiveness over time as they are blocked by spam filters.... This suggests the rock-phish gang are motivated to purchase new domains even when registrars are slow to take action."

Effect of and Response to Malicious Registrations

The highest-scoring TLDs almost invariably suffered from the systematic registration of domain names by phishers:

#1: .LY (Libya. Score 271.0; 84 phishing domains out of 3,100 domains in the registry.) Most of the .LY phishing domains were maliciously registered (in the BIZ.LY zone), and contained brand names.

#2: .MN (Mongolia. Score 182.2; 93 phishing domains out of 4,984 domains in the registry.) Of the 93 domains, 80 were methodically registered by one phisher.

#3: .HK (Hong Kong. Score 113.2; 1,717 phishing domains out of 150,799 domains in the registry.) Phishers (including the Rock Phish gang) systematically exploited weaknesses in the .HK's registry's anti-abuse capabilities. This story illustrates how phishers "discover" new TLDs that are useful for their purposes, and exploit them for as long as it is effective. The good news in this story is that the registry operator, HKDNR, developed effective responses and brought the attacks to an end. HKDNR has been sharing its experiences with other registries and groups, and here are some of the more interesting highlights:

a. HKDNR is not only the registry operator, but is also the sole retail registrar for .HK domain names. Thus, if someone wants to register a .HK domain name, he or she must use the HKDNR Web site.

[10] Tyler Moore and Richard Clayton, Computer Laboratory, University of Cambridge: "Examining the Impact of Website Take-down on Phishing," http://www.cl.cam.ac.uk/~rnc1/ecrime07.pdf

b. HKDNR has very robust functionality for managing DNS, including rapid zone updates for new domains and changes. This is particularly appealing to phishers and other criminals who wish to use techniques like fast flux and botnets to host sites or send spam from.
c. It is quick and easy to register domain names using the HKDNR Web site. Before the attacks began, the process appears to have been relatively easy to automate. As a countermeasure, HKDNR installed more checks to spot abusive automated registration behaviors.
d. Before the attacks began, HKDNR had a domain suspension policy that did not take into account the use of .HK domain names for criminal activities, and did not provide for domain suspensions. Instead, HKDNR relied on a traditional dispute policy that assumed that domain use problems would be based on trademark or infringement claims. If a domain seemed to be used for criminal activity, HKDRN was obligated to report the issue to the Hong Kong police department, which would then take two weeks or more to investigate and issue an order to suspend the site. During the attacks, HKDNR revised its policy in consultation with the Honk Kong police department, APWG members, and other interested parties. HKDNR now has a fast-track process so that phishing domains are typically suspended in half a day or less.
e. Before the attacks, credit card verification on the HKDNR site was not performed in real-time with the latest PCI (Payment Card Industry) techniques for detecting fraud. This allowed criminals to easily use stolen credit card information to register domains. In reaction to these attacks, HKDNR strengthened its card-processing system to use more fraud detection techniques.

#7: .TH (Thailand. Score: 58.1; 171 phishing domains out of 33,000 domains in the registry). 101 of the phishing domains were registered suspiciously under the AC.TH zone.

#8: .LI (Liechtenstein. Score: 44.7; with 221 phishing domains out of 50,100 domains in the registry.) It appears that at least 201 of the 221 domains were registered maliciously, many by the Rock Phish gang, and were used to concurrently target companies of various types around the world.

At the low end of the scale, .CN is a notable case. .CN had a phishing score of only 2.2, with 1,853 phishing domains out of 8,459,174 domains in the registry. However, at least 1,504 of those domains (81%) appear to have been maliciously registered by phishers. In 2008, it appears that .CN names are being registered heavily by spammers and phishers, and .CN's score in 2008 may rise significantly.

Conclusion

As always, 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 adoption of subdomain services, evasion and spoofing techniques, and their systematic exploitation of vulnerable registrars and registries. We hope this study will spur further research on these and related topics.

The number of domain names used for phishing in 2007 was upwards of 52,000. This was a miniscule percentage of the approximately 153 million total domain names in existence, but the phishing resulted in huge financial losses for Internet users and the targeted brands. We have noted some of the problems associated with detecting and mitigating phishing in this ocean of domain names. 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 both can mitigate malicious domain name registrations quickly, thereby reducing phishing up-times and reducing the options available to phishers.

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, and .IN (India). Greg oversees 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), and has advised the Government of India regarding domain and related Internet policies. 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. He 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 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.

Appendix A: TLD Phishing Scores

We were able to obtain the number of domains in the below 105 TLD registries. See Appendix B for a list of all TLDs that contained phishing domains.

Rank	TLD	TLD Location	Domains in registry in November 2007	Domain names used for phishing in 2007	Score: Phish per 10,000 domains
1	.ly	Libya	3,100	84	**271.0**
2	.mn	Mongolia	5,087	93	**182.8**
3	.hk	Hong Kong	150,799	1,707	**113.2**
4	.edu	U.S. education	6,997	67	**95.8**
5	.al	Albania	250	2	**80.0**
6	.md	Moldova	2,200	15	**68.2**
7	.th	Thailand	33,000	171	**51.8**
8	.li	Liechtenstein	50,100	221	**44.1**
9	.hn	Honduras	3,820	16	**41.9**
10	.co	Colombia	20,524	65	**31.7**
11	.bo	Bolivia	3,705	11	**29.7**
12	.cx	Christmas Island	4,387	13	**29.6**
13	.tc	Turks and Caicos	9,000	20	**22.2**
14	.vg	British Virgin Islands	7,405	15	**20.3**
15	.ba	Bosnia and Herzegovina	6,606	13	**19.7**
16	.ec	Ecuador	14,941	29	**19.4**
17	.pe	Peru	17,859	33	**18.5**
18	.bg	Bulgaria	7,500	13	**17.3**
19	.py	Paraguay	6,501	10	**15.4**
20	.gt	Guatemala	6,262	9	**14.4**
21	.am	Armenia	8,570	12	**14.0**
22	.cu	Cuba	1,455	2	**13.7**
23	.ro	Romania	242,484	316	**13.0**
24	.cl	Chile	195,513	222	**11.4**

25	.bz	Belize	42,360	48	**11.3**
26	.np	Nepal	11,016	12	**10.9**
27	.tw	Taiwan	341,462	361	**10.6**
28	.lv	Latvia	28,900	30	**10.4**
29	.lt	Lithuania	64,554	65	**10.1**
30	.ee	Estonia	50,000	47	**9.4**
31	.su	Soviet Union	19,431	17	**8.7**
32	.cz	Czech Republic	347,989	286	**8.2**
33	.mx	Mexico	230,177	189	**8.2**
34	.is	Iceland	20,000	16	**8.0**
35	.uy	Uruguay	13,936	11	**7.9**
36	.pl	Poland	753,520	581	**7.7**
37	.sv	El Salvador	4,184	3	**7.2**
38	.sk	Slovakia	150,601	107	**7.1**
39	.ni	Nicaragua	4,254	3	**7.1**
40	.ve	Venezuela	53,704	36	**6.7**
41	.pa	Panama	4,488	3	**6.7**
42	.yu	Yugoslavia	46,279	30	**6.5**
43	.sa	Saudi Arabia	12,478	8	**6.4**
44	gi	Gibraltar	1,602	1	**6.2**
45	.ke	Kenya	8,011	5	**6.2**
46	.ru	Russia	1,104,572	684	**6.2**
47	.at	Austria	722,193	415	**5.7**
48	do	Dominican Republic	10,873	6	**5.5**
49	.tr	Turkey	142,646	73	**5.1**
50	.in	India	331,495	168	**5.1**
51	.hu	Hungary	350,000	173	**4.9**
52	.us	United States	1,362,805	661	**4.9**
53	.sg	Singapore	87,086	41	**4.7**
54	.be	Belgium	726,000	340	**4.7**
55	.cat	sponsored TLD	25,885	12	**4.6**
56	.ch	Switzerland	1,036,000	470	**4.5**
57	.br	Brazil	1,262,967	563	**4.5**
58	.gr	Greece	202,000	88	**4.4**
59	.cr	Costa Rica	6,905	3	**4.3**
60	.pt	Portugal	184,596	80	**4.3**
61	.ua	Ukraine	311,822	135	**4.3**

62	.nz	New Zealand	311,198	134	**4.3**
63	.kr	Korea	932,841	394	**4.2**
64	.my	Malaysia	98,000	41	**4.2**
65	.org	generic TLD	6,412,064	2,627	**4.1**
66	.hr	Croatia	51,432	21	**4.1**
67	.si	Slovenia	50,312	20	**4.0**
68	.biz	generic TLD	1,944,453	764	**3.9**
69	.il	Israel	112,500	43	**3.8**
70	.net	generic TLD	10,581,849	3,973	**3.8**
71	.jp	Japan	972,584	359	**3.7**
72	.aero	sponsored TLD	5,430	2	**3.7**
73	.mu	Mauritius	5,500	2	**3.6**
74	.ir	Iran	72,906	26	**3.6**
75	.vc	Saint Vincent and the Grenadines	5,662	2	**3.5**
76	.za	South Africa	359,518	126	**3.5**
77	.ma	Morocco	25,873	9	**3.5**
78	.es	Spain	770,984	263	**3.4**
79	.com	generic TLD	70,698,420	23,860	**3.4**
80	.au	Australia	985,458	314	**3.2**
81	.fr	France	969,864	307	**3.2**
82	.ca	Canada	935,000	286	**3.1**
83	.lu	Luxembourg	34,000	10	**2.9**
84	.dk	Denmark	862,000	239	**2.8**
85	.nl	Netherlands	2,661,308	737	**2.8**
86	.ie	Ireland	90,710	25	**2.8**
87	.it	Italy	1,467,221	401	**2.7**
88	.info	generic TLD	4,954,266	1,295	**2.6**
89	.no	Norway	357,722	92	**2.6**
90	.cy	Cyprus	8,229	2	**2.4**
91	.fi	Finland	165,000	38	**2.3**
92	.ag	Antigua and Barbuda	13,507	3	**2.2**
93	.vn	Vietnam	54,739	12	**2.2**
94	.cn	China	8,459,174	1,853	**2.2**
95	.ws	Samoa	522,221	114	**2.2**
96	.name	sponsored TLD	265,638	55	**2.1**

97	.se	Sweden	685,000	127	**1.9**
98	.ar	Argentina	1,451,727	230	**1.6**
99	.de	Germany	11,524,091	1,798	**1.6**
100	.uk	United Kingdom	6,445,465	992	**1.5**
101	.im	Isle of Man	8,500	1	**1.2**
102	.eu	European Union	2,671,846	197	**0.7**
103	.mobi	sponsored TLD	761,549	48	**0.6**
104	.dm	Dominica	19,469	1	**0.5**
105	.travel	sponsored TLD	28,665	1	**0.3**
	TOTALS		**150,698,751**	**50,774**	

Appendix B: Phishing by TLD

TLD	TLD Location	Domains in registry in November 2007	Domain names used for phishing in 2007	Score: Phishing domains per 10,000
.ac	Ascension Island		5	
.ae	United Arab Emirates		5	
.aero	sponsored TLD	5,430	2	**3.7**
.ag	Antigua and Barbuda	13,507	3	**2.2**
.ai	Anguilla		4	
.al	Albania	250	2	**80.0**
.am	Armenia	8,570	12	**14.0**
.ar	Argentina	1,451,727	230	**1.6**
.as	American Samoa		7	
.at	Austria	722,193	415	**5.7**
.au	Australia	985,458	314	**3.2**
.az	Azerbaijan		1	
.ba	Bosnia and Herzegovina	6,606	13	**19.7**
.bd	Bangladesh		6	
.be	Belgium	726,000	340	**4.7**
.bf	Burkina Faso		2	
.bg	Bulgaria	7,500	13	**17.3**
.bi	Burundi		1	
.biz	generic TLD	1,944,453	764	**3.9**
.bm	Bermuda		1	
.bn	Brunei Darussalam		4	
.bo	Bolivia	3,705	11	**29.7**
.br	Brazil	1,262,967	563	**4.5**
.bs	Bahamas		1	
.by	Belarus		12	
.bz	Belize	42,360	48	**11.3**
.ca	Canada	935,000	286	**3.1**
.cat	sponsored TLD	25,885	12	**4.6**

.cc	Cocos (Keeling) Islands		161	
.cd	Democratic Republic of the Congo		6	
.ch	Switzerland	1,036,000	470	**4.5**
.ci	Cote D'Ivoire (Ivory Coast)		3	
.cl	Chile	195,513	222	**11.4**
.cn	China	8,459,174	1,853	**2.2**
.co	Colombia	20,524	65	**31.7**
.com	generic TLD	70,698,420	23,860	**3.4**
.coop	sponsored TLD		3	
.cr	Costa Rica	6,905	3	**4.3**
.cu	Cuba	1,455	2	**13.7**
.cx	Christmas Island	4,387	13	**29.6**
.cy	Cyprus	8,229	2	**2.4**
.cz	Czech Republic	347,989	286	**8.2**
.de	Germany	11,524,091	1,798	**1.6**
.dk	Denmark	862,000	239	**2.8**
.dm	Dominica	19,469	1	**0.5**
.do	Dominican Republic	10,873	6	**5.5**
.ec	Ecuador	14,941	29	**19.4**
.edu	U.S. education	6,997	67	**95.8**
.ee	Estonia	50,000	47	**9.4**
.eg	Egypt		5	
.es	Spain	770,984	263	**3.4**
.et	Ethiopia		1	
.eu	European Union	2,671,846	197	**0.7**
.fi	Finland	165,000	38	**2.3**
.fm	Federated States of Micronesia		12	
.fo	Faroe Islands		1	
.fr	France	969,864	307	**3.2**
.gd	Grenada		2	
.ge	Georgia		5	
.gg	Guernsey		2	
.gh	Ghana		6	
.gi	Gibraltar	1,602	1	**6.2**
.gm	Gambia		3	

.gov	U.S. government		2	
.gp	Guadeloupe		1	
.gr	Greece	202,000	88	**4.4**
.gs	S. Georgia and S. Sandwich Islands		7	
.gt	Guatemala	6,262	9	**14.4**
.hk	Hong Kong	150,799	1,707	**113.2**
.hm	Heard and McDonald Islands		3	
.hn	Honduras	3,820	16	**41.9**
.hr	Croatia	51,432	21	**4.1**
.hu	Hungary	350,000	173	**4.9**
.id	Indonesia		60	
.ie	Ireland	90,710	25	**2.8**
.il	Israel	112,500	43	**3.8**
.im	Isle of Man	8,500	1	**1.2**
.in	India	331,495	168	**5.1**
.info	generic TLD	4,954,266	1,295	**2.6**
.int	sponsored TLD		1	
.io	British Indian Ocean Territory		28	
.ir	Iran	72,906	26	**3.6**
.is	Iceland	20,000	16	**8.0**
.it	Italy	1,467,221	401	**2.7**
.jo	Jordan		5	
.jp	Japan	972,584	359	**3.7**
.ke	Kenya	8,011	5	**6.2**
.kg	Kyrgyzstan		23	
.kh	Cambodia		1	
.kr	Korea	932,841	394	**4.2**
.kw	Kuwait		1	
.kz	Kazakhstan		15	
.la	Laos		16	
.li	Liechtenstein	50,100	221	**44.1**
.lk	Sri Lanka		11	
.lt	Lithuania	64,554	65	**10.1**
.lu	Luxembourg	34,000	10	**2.9**

.lv	Latvia	28,900	30	**10.4**
.ly	Libya	3,100	84	**271.0**
.ma	Morocco	25,873	9	**3.5**
.md	Moldova	2,200	15	**68.2**
.mg	Madagascar		5	
.mk	Macedonia		7	
.mn	Mongolia	5,087	93	**182.8**
.mo	Macao		4	
.mobi	sponsored TLD	761,549	48	**0.6**
.ms	Montserrat		23	
.mt	Malta		2	
.mu	Mauritius	5,500	2	**3.6**
.mx	Mexico	230,177	189	**8.2**
.my	Malaysia	98,000	41	**4.2**
.mz	Mozambique		2	
.na	Namibia		2	
.name	sponsored TLD	265,638	55	**2.1**
.ne	Niger		1	
.net	generic TLD	10,581,849	3,973	**3.8**
.nf	Norfolk Island		1	
.ng	Nigeria		2	
.ni	Nicaragua	4,254	3	**7.1**
.nl	Netherlands	2,661,308	737	**2.8**
.no	Norway	357,722	92	**2.6**
.np	Nepal	11,016	12	**10.9**
.nr	Nauru		12	
.nu	Niue		89	
.nz	New Zealand	311,198	134	**4.3**
.org	generic TLD	6,412,064	2,627	**4.1**
.pa	Panama	4,488	3	**6.7**
.pe	Peru	17,859	33	**18.5**
.ph	Philippines		196	
.pk	Pakistan		22	
.pl	Poland	753,520	581	**7.7**
.pn	Pitcairn		5	
.ps	Palestinian Territory		11	
.pt	Portugal	184,596	80	**4.3**
.py	Paraguay	6,501	10	**15.4**

.ro	Romania	242,484	316	**13.0**
.ru	Russia	1,104,572	684	**6.2**
.rw	Rwanda		1	
.sa	Saudi Arabia	12,478	8	**6.4**
.sd	Sudan		2	
.se	Sweden	685,000	127	**1.9**
.sg	Singapore	87,086	41	**4.7**
.sh	Saint Helena		9	
.si	Slovenia	50,312	20	**4.0**
.sk	Slovakia	150,601	107	**7.1**
.sn	Senegal		2	
.st	Sao Tome and Principe		49	
.su	Soviet Union	19,431	17	**8.7**
.sv	El Salvador	4,184	3	**7.2**
.tc	Turks and Caicos	9,000	20	**22.2**
.tf	French Southern Territories		4	
.tg	Togo		1	
.th	Thailand	33,000	171	**51.8**
.tj	Tajikistan		2	
.tk	Tokelau		102	
.tl	Timor-Leste		2	
.tm	Turkmenistan		1	
.tn	Tunisia		4	
.to	Tonga		29	
.tp	Portuguese Timor		4	
.tr	Turkey	142,646	73	**5.1**
.travel	sponsored TLD	28,665	1	**0.3**
.tt	Trinidad and Tobago		3	
.tv	Tuvalu		144	
.tw	Taiwan	341,462	361	**10.6**
.tz	Tanzania		5	
.ua	Ukraine	311,822	135	**4.3**
.ug	Uganda		11	
.uk	United Kingdom	6,445,465	992	**1.5**
.us	United States	1,362,805	661	**4.9**
.uy	Uruguay	13,936	11	**7.9**
.uz	Uzbekistan		3	

.vc	Saint Vincent and the Grenadines	5,662	2	**3.5**
.ve	Venezuela	53,704	36	**6.7**
.vg	British Virgin Islands	7,405	15	**20.3**
.vi	U.S. Virgin islands		1	
.vn	Vietnam	54,739	12	**2.2**
.vu	Vanuatu		12	
.ws	Samoa	522,221	114	**2.2**
.yu	Yugoslavia	46,279	30	**6.5**
.za	South Africa	359,518	126	**3.5**
.zw	Zimbabwe		7	
TOTALS		**150,698,751**	**51,989**	

#