Potential for Phishing in Sensitive-String Top-Level Domains

A Study for the ICANN Board of Directors New

gTLD Program Committee

ILLUMINTEL

21 May, 2015

Table of Contents

Executive Summary ... 3

Conclusions ... 3

Data Set and Methodology ... 5

How Phishing Works ... 7

Spam Lures ... 7

Domain Name Selection by Phishers ... 10

Compromised Domains ... 12

Malicious Registrations ... 14

Subdomain Registries ... 16

IP Addresses ... 17

Domain Choice and URL Construction by Phishers ... 19

Internationalized Domain Names ... 21

Phishing Success Rates ... 23

Browser Warnings ... 23

Click Rates ... 24

Identity Theft Success Rates ... 24

Phishing Uptimes ... 25

New gTLD Analysis ... 26

Prevention and Mitigation ... 29

Related Scams ... 31

Acknowledgments ... 32

About the Author ... 32

Appendix: Phishing Statistics by TLD, 2012-2014 ... 33

Executive Summary

In April 2013, ICANN's Governmental Advisory Committee (GAC) advised the ICANN Board that specific safeguards should apply to various categories of new TLD strings related to regulated activities. Examples included strings that described sectors "such as those subject to national regulation (such as .bank, .pharmacy) or those that describe or are targeted to a population or industry that is vulnerable to online fraud or abuse." A concern was that these strings may convey meaning or sense of trust to consumers, and consumers may be under the impression that these spaces are somehow regulated or limited to qualified registrants.

On 5 February 2014, the ICANN Board's New gTLD Program Committee (NGPC) adopted an implementation framework [PDF, 61 KB] for GAC Category 1 Safeguard Advice. The implementation framework classifies each Category 1 string as requiring one of three levels of safeguards. The adoption of the implementation framework has allowed applicants subject to GAC Category 1 Advice to proceed in the New gTLD Program once other eligibility criteria have been met. The implementation framework requires safeguards to be added to Specification 11 of the Registry Agreement as public interest commitments. For these applications, these safeguards are mandatory requirements.[1]

The issue of consumer harm has remained of interest to some members of the ICANN community. At the request of the NGPC, this report explores if and how sensitive-string TLDs can be exploited by phishers and related scammers, and therefore in that respect whether sensitive-string TLDs pose a higher or lower risk of increased harm or abuse to Internet users.

This intention of this report is to provide informative, fact-based analysis. This report does not study other types of abuse (such as malware); and does not study other aspects of the sensitive strings topic, such as the security risks of Web sites (registrants) that hold sensitive data in regulated fields such as finance and healthcare, the risk of intellectual property infringement, or the effects of registrant validation methods.

Conclusions

A. Phishers and other scammers can successfully leverage the meaning associated with a TLD string if the potential victim:

1. can see the URL in the phishing email or in their web browser window, *and*
2. can correctly identify the domain name and TLD string in that URL, *and*

[1] Source: http://newgtlds.icann.org/en/applicants/gac-advice/cat1-safeguards

3. places trust in the TLD string in question.

Regarding #1: Phishing is often designed to hide the real destination domain name from the user, rather than to emphasize the domain. (Pages 8-9)

Regarding #2: Many Internet users do not have the knowledge needed to recognize the domain name (and therefore the TLD string) in a URL. Usually the domain registrations made by phishers consist of nonsense strings that have no semantic meaning whatsoever. (Pages 19-21, 24)

Regarding #2 and #3: It is unknown how many Internet users are aware of the GAC-designated sensitive strings and recognize them as officially delegated TLDs, although the number is certainly small at this time. It is difficult to quantify how many Internet users may assign trust to a specific TLD string, or to what extent. Finally, many Internet users ignore phishing warnings. (Pages 23-24)

B. A phisher does not need to register a domain in a sensitive string TLD in order to get the semantic benefit. He can simply create a URL string that *appears* to be in the sensitive TLD. A phisher can use any domain name in any TLD to spoof any TLD string, including a sensitive-string TLD. (Pages 19-21)

C. Most phishing takes place on compromised domains, where the phisher has broken into the registrant's web hosting. As a result, most top-level domains experience phishing in them—even TLDs that have registration restrictions. Because of phishing on compromised domains, the ICANN community should expect that many new gTLDs will experience phishing over time. (Pages 13-14, 26)

D. The scale of phishing is very small compared to the number of domains in the world. Phishers registered less than 100,000 domains over the last three years in all TLDs worldwide, mostly concentrated in certain legacy gTLDs and ccTLDs. (Pages 10-11, 19-20)

E. Historically, phishers have had an almost infinite supply of resources they can use for phishing. These include domains available for registration in the world's legacy gTLD and ccTLD registries (pages 14-16); domains on vulnerable web hosting; free subdomains (pages 16-17); and IP-based phishing (pages 17-18). So the current expansion of the gTLD domain space will probably not increase the *total amount* of phishing in the world. *Instead, it will create new or different locations where phishing occurs* in the DNS, since cyber-criminals move from TLD to TLD over time. (Pages 15-16, 22, 26)

F. Phishing can be deterred by measures including registration restrictions, pricing strategies, and active mitigation. (Pages 16, 29-30) Detection and mitigation efforts determine how long phishing sites stay up, and therefore how much harm phishing attacks cause. (Page 25) So, vital issues are *where* in the DNS phishing takes place, *how long* phishing attacks stay up, and what hosting providers, registries, and registrars do to prevent it from happening or stop it once it does happen.

Data Set and Methodology

The primary domain name and phishing data in this report comes from the Anti-Phishing Working Group's *Global Phishing Survey* series. The reports are the security industry's authoritative studies of phishing metrics and the use of domain names for phishing. Authored by APWG contributors Greg Aaron of Illumintel Inc. and Rod Rasmussen of Internet Identity (IID), the series has been published semi-annually since 2008, with each report covering a half-year period. The complete set of reports is available at: http://www.apwg.org/resources/apwg-reports/whitepapers

This report for ICANN concentrates on the phishing that was documented in the three years between 1 January 2012 to 31 December 2014.

The *Global Phishing Survey* data was compiled from several sources. The largest source was the Anti-Phishing Working Group itself, which operates the security industry's major repository of phishing and e-mail fraud activity. This includes the contents of the APWG's real-time feed of live phishing URLs, which is provided by APWG members, many of whom are either phishing targets or security companies that assist phishing targets and provide products and data to protect Internet users. The APWG data was supplemented with data from IID (which operates honeypots and receives various intelligence feeds), phishing feeds such as PhishTank, data kindly provided by CNNIC and the Anti-Phishing Alliance of China (which acts as the official clearinghouse for phishing reporting and mitigation in China), and other private sources.

The *Global Phishing Survey* methodology is designed to collect as many valid phishing URLs as possible, to confirm the reliability of that data, and to discern meaning from the data. A prime goal has been to learn how many unique phishing attacks and domain names are involved worldwide. Millions of phishing URLs are reported by our sources each year, but we determine that the number of unique phishing attacks and the number of domain names used to host them is much smaller.[2]

The phishing pages were verified during the course of the attacks -- either by an automated check performed by Internet Identity that confirmed the presence and characteristics of the destination phishing page (a check that often also captured a screenshot and page source code), and/or by a human analyst at IID or CNNIC/APAC. All unconfirmed phishing attacks were discarded, and so the data set does not contain false-positives. The data set certainly does not capture all of the phishing attacks that

[2] This is due to several factors: A) Some phishing involves customized attacks that incorporate unique numbers in the URLs, often to track targeted victims, or to defeat spam filters. A singe phishing attack can therefore manifest as thousands to millions of individual URLs, which all lead to essentially one phishing site. Counting all URLs would therefore inflate the importance and impact of some phishing campaigns. Our counting method takes URL "paths" into account and de-duplicates in order to count unique attacks, and this method has remained consistent across our reports. Other observers de-duplicate and count differently. B) Phishers often use one domain name to host simultaneous attacks against different targets. Some phishers place several different phishing attacks on each domain name they register. C) A phishing site may have multiple pages, each of which may be reported.

take place around the world, but we believe that our data set is one of the largest and most carefully constructed available.

The data set consists of phishing aimed at the general Internet-using public. It does not contain spear-phishing data. Spear phishing is directed against specific individuals or entities into order to gain access to that organization's systems. Spear-phishing is often not reported, and is not aimed at a wider Internet-using population.

Our definitions are:

- Attack: a phishing site that targets a specific brand or entity. A single domain name can host several different phishing attacks against different banks, for example.
- Domain name: a second-level domain name, or a third-level domain name if the relevant registry offers third-level registrations. (An example is the .UK registry, which offers both second-level registrations and third-level registrations in zones such as CO.UK and ORG.UK.) In other words, domain names that can be registered in a TLD registry.

For this report we also examined the 300 most recent domains listed at Artists Against 419 (aa419.org), a large repository of fraud sites, mainly advance fee fraud sites (see page 31).

How Phishing Works

Phishing is one of the major methods for perpetrating fraud and identity theft on the Internet. It is a "social engineering" attack in that it preys on the victim's trust and inexperience.

Specifically, phishing is an attempt to acquire sensitive information such as usernames, passwords, and credit card details by masquerading as a trustworthy entity. Phishers create web sites that masquerade as sites that an Internet user trusts, such as the site of this bank or favorite social media site. Most often, phishers lure Internet users to these bogus sites via legitimate-looking emails (spam), which are also made to appear legitimate. These mails usually contain an incentive to click to the bogus site and enter personal information, such as a statement that the user's account information needs to be updated.

It is easy to set up a phishing attack. Phishing "kits" are available on the underground market; they allow anyone with a modicum of computer skill to create a professional-looking data-stealing campaign. Phishing kits contain Web page templates, data extraction tools that allow the phisher to collect the harvested data, and spamming tools. Sophisticated phishers automate their work, allowing them to launch new phishing sites easily, keeping ahead of the mitigation attempts of security responders. Automation lowers the phishers' costs, and makes phishing lucrative even when only a tiny percentage of potential victims fall for the scam.

The losses due to phishing cannot be reliably quantified. Victim companies such as banks do not release the amount of money involved, many victims do not report their losses, some phishing attacks do not lead to immediate financial losses, and there is no worldwide clearinghouse that collects consumer complaints. Most loss estimates run from the hundreds of millions to billions of U.S. dollars per year.[3] Phishing also imposes substantial indirect costs, such as the effort that it takes for victims to report the crime and recover from the damage to their accounts and credit ratings.

According to our *Global Phishing Survey* statistics, the number of phishing attacks has doubled since 2008. Clearly phishing is a worthwhile endeavor for criminals.

Spam Lures

It is trivially easy to forge the sender ("from") address in an email[4], and phishers usually do so in order to make their emails look like they are coming from the trusted institution. Many phishing emails are in HTML format, and will often feature a link or "click here" button. These are usually configured to hide

[3] At the high end, EMC/RSA estimated US$5.8 billion in losses in 2014: http://www.emc.com/collateral/fraud-report/rsa-online-fraud-report-012014.pdf

[4] See http://en.wikipedia.org/wiki/Email_spoofing and http://www.huffingtonpost.com/jason-p-stadtlander/email-spoofing-explained-_1_b_6477672.html

the true URL that the user will be directed to. If the user hovers the cursor over a link, his email client or web browser will preview the real link URL. But even this helpful feature can be overridden by the more resourceful phishers.

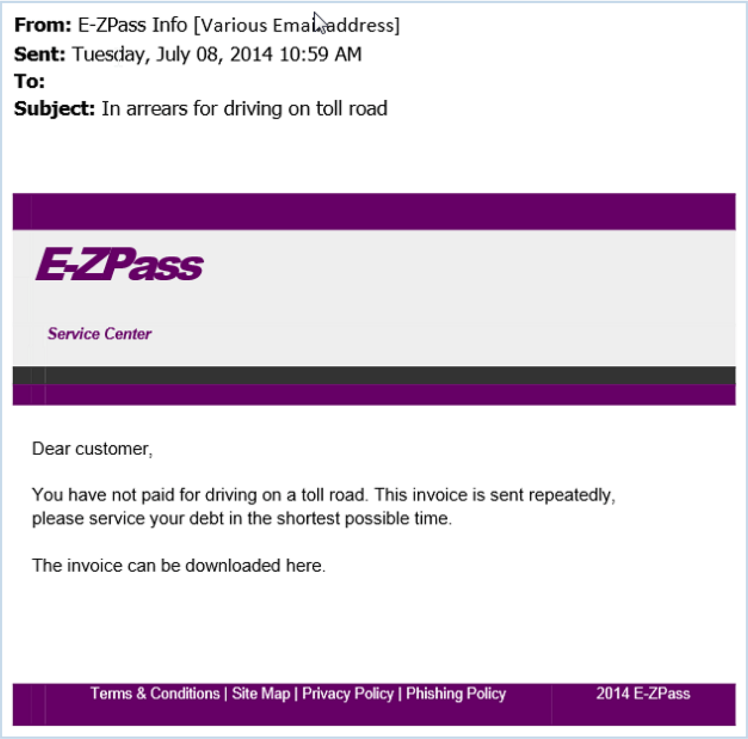

From: E-ZPass Info [Various Email address]
Sent: Tuesday, July 08, 2014 10:59 AM
To:
Subject: In arrears for driving on toll road

Dear customer,

You have not paid for driving on a toll road. This invoice is sent repeatedly, please service your debt in the shortest possible time.

The invoice can be downloaded here.

Terms & Conditions | Site Map | Privacy Policy | Phishing Policy
2014 E-ZPass

Above: in this phishing email, the "From" email address is spoofed, and was not sent from E-ZPass

Phishers mainly use compromised domains (see pages 12-14), and those registered domain strings have no relation to the trusted site that the phisher is spoofing. In such cases, the phishers hide the real URL/domain name because some potential victims may be able to discern that the advertised domain is not that of the real, trusted site.

Sometimes the links that phishers place in their email lures are to redirects. This method also disguises the location of the end phishing page from the Internet user. These links are usually either URLs at URL shortener services, or links on an innocent web site that has been hacked into by the phisher. When the user clicks on the link in the email, the user is sent to the intermediary URL/domain, which then automatically redirects the user onwards to the final destination. And since the intermediary domain may be a trusted one, this method helps the phisher get his email past anti-spam filtering.

Some phishers even use methods that alter what is displayed in the user's browser address bar, although these methods are rarer. This is done either by placing a picture of a legitimate URL over the address bar, or by closing the original bar and opening up a new one.

A minority of phishing sites also incorporate malware of various types, which may give the phisher another way to capture the user's sensitive information.

Conclusion: **Phishing emails are often designed to hide the real destination domain name from the user, rather than to emphasize the domain.**

Domain Name Selection by Phishers

To mount a phishing attack, a phisher needs to set up a phishing site accessible to users on the Web. Phishers do so using the following types of resources:

1. *Compromised domains.* These are domains where the phisher has broken into the web hosting of an innocent registrant.
2. *Malicious registrations.* These are domains registered by phishers, for phishing.
3. *Subdomain resellers.* These services allow people to create registrations at the third level beneath a second-level domain name that the service provider owns.
4. *IP addresses.* The URLs of these phishing attacks contain IP addresses rather than domain names.

These different categories are important because they present different mitigation options for responders, and offer insights into how phishers commit their crimes. Below we examine these four categories in more detail.

Following are the numbers from the *APWG Global Phishing Surveys* covering the half-year periods from 1 January 2012 through 31 December 2014:

	1H2012	2H2012	1H2013	2H2013	1H2014	2H2014	Totals 2012-2014
Domains used for phishing[5]	64,204	89,748	53,685	82,163	87,901	95,321	**473,022**
Phishing attacks	93,462	123,476	72,758	115,565	123,741	123,973	**652,975**
Compromised domains	55,492	83,029	40,448	58,398	64,489	67,281	**369,137**
Maliciously registered domains	7,712	5,833	12,173	22,831	22,679	27,253	**98,481**
IP-based phish (unique IPs)	1,864	1,841	1,626	837	2,317	3,095	**11,580**

The Appendix of this report contains statistics for all TLDs world-wide, including TLD size, number of phishing attacks, scoring, and number of malicious registrations.

To put the above numbers in perspective:

[5] Domains used for phishing = malicious domains + compromised domains + domains at URL shortening services and subdomain resellers.

- As of December 2014, there were approximately 287.3 million domain names in the world's registries.[6]
- Registrations in the world's registries grew by 16.9 million, or 6.2 percent, in the year Q4 2013 to Q4 2014.[7]
- During the period 2012-2014, there were many more than 287.3 million domain names in existence. Many domains expired and were purged from registries during that period; some were unique and some were then re-registered. The annual renewal rate for domains in the world's registries may have averaged around 75%.[8] If a domain expired and was deleted from its registry and was then re-registered we can consider that a "new" registration or a new "domain create." Based on a 75% renewal rate, there may have been around 84.6 million domains created in 2014.[9]
- By this logic there would have been in excess of 400 million total domain names in and out of existence at some point in the three years 2012 through 2014.

[6] Based on our collection of statistics from ICANN registry reports, ccTLD registries, and nTLDSTATS.COM. VeriSign and ZookNIC estimate 288 million domains: http://www.verisigninc.com/assets/domain-name-report-march2015.pdf

[7] VeriSign Domain Name Industry Brief, March 2015: http://www.verisigninc.com/assets/domain-name-report-march2015.pdf

[8] The recent renewal rate for .COM has been around 72%, with some gTLDs reporting lower renewal rates and some large ccTLDs reporting higher renewal rates. Statistics provided by Verisign and ZookNIC, at; http://www.verisigninc.com/assets/domain-name-report-march2015.pdf

[9] 270.4 million domains in the world at the beginning of 2014, renewing at a 75% rate equals 67.7 million creates, plus 16.9 million domains of net growth, arriving at 287.3 million domains at the end of 2014.

Compromised Domains

Most phishing attacks at hosted on compromised domains. These are domains where the web hosting has been hacked into by a phisher. In recent years, phishers have especially preyed on web sites that are running on out-of-date versions of WordPress and Joomla. Phishers often use automated tools that scan the Internet for vulnerable hosting. Password compromise can also enable break-ins.

The domain's registrant typically has no idea that the break-in has happened. The phisher inserts the phishing page in a new subdomain or subdirectory of the web site, where is will not be encountered by the site's usual visitors.

Below is a typical example. The domain kendenup.com was registered in 2002, and is the web site of a hotel in Perth, Australia. The site is hosted at ISP Omniconnect in Melbourne. The home page displays the innocent registrant's site:

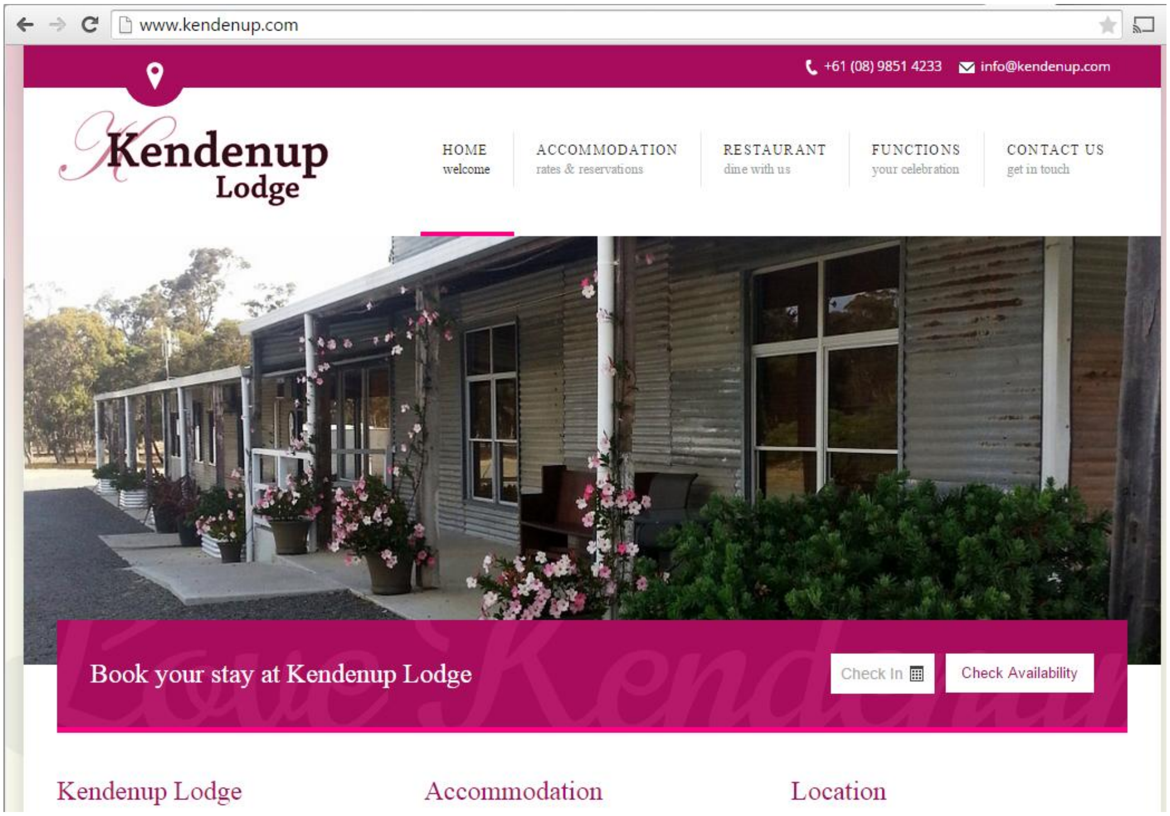

Around 19 April 2015, the hotel's web hosting was broken into a phisher, who installed the phishing page below, which imitates Google Docs. The phishing attack page was placed in a subdirectory, at: http://kendenup.com/secure/psp/note/

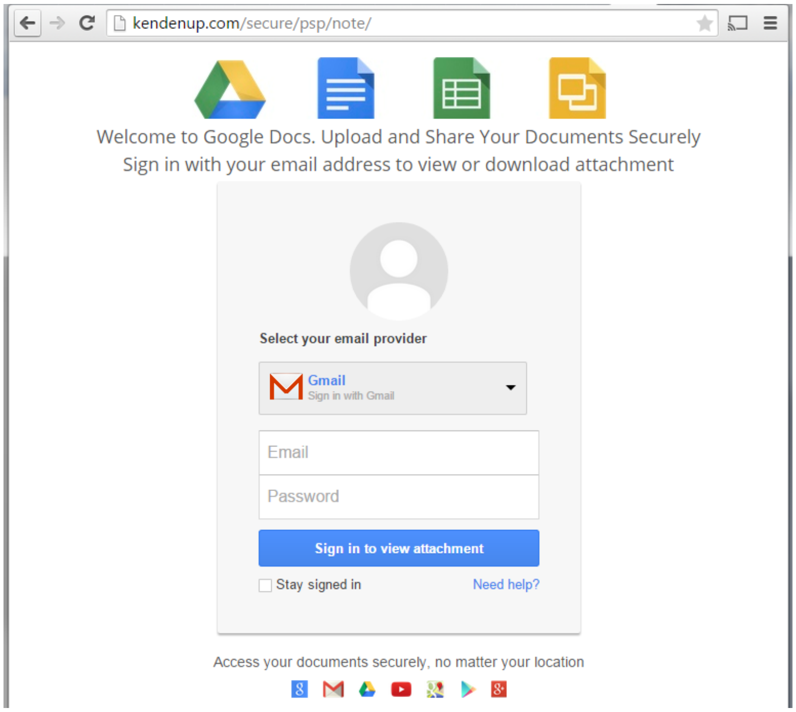

kendenup.com/secure/psp/note/

Welcome to Google Docs. Upload and Share Your Documents Securely
Sign in with your email address to view or download attachment

Select your email provider

Access your documents securely, no matter your location

Above: phishing page at http://kendenup.com/secure/psp/note/

Phishers use compromised domains for three reasons:

1. The compromised domain will have a *good reputation* – it will have no history of abuse, and it will have a domain registration date that is not recent and thus scores better on reputational checks. Such domains are advantageous to use in phishing emails because the domains are legitimate and are less likely to be blocklisted and filtered out by anti-spam programs.
2. *Mitigating these phish is harder*. Such domains are not candidates for suspension by the registrar or registry, reducing the options for mitigation. If the domain name is suspended, the entire legitimate web site (and associated domain-dependent services such as email) will also go down. This harms the legitimate registrant and all the users of and visitors to the domain. Instead, the preferred mitigation option is to have the hosting provider remove the phishing pages at the server, and then patch the server to prevent further break-ins.
3. Using a compromised site helps the phisher *avoid detection*. The phisher does not need to purchase a domain name or hosting, which will leave a transaction trail, and he often remains safe in another legal jurisdiction.

APWG studies consistently find that the United States is the top country where phishing sites are hosted, simply because the United States has so many hosting providers and web sites/domains hosted within its borders.[10] The .COM TLD contains 41.3% of the domains in the world, and 58% of the phishing domains in our 2H2014 data set – an expected concentration of compromised domains due to .COM's established size and ubiquity, plus a large number of maliciously registered domains (perhaps chosen by phishers due to .COM's ubiquity).

APWG studies also indicate how hosting security in some places is weaker than in others; in 2014 hosting in Latin America was disproportionately vulnerable.[11] In another example from our *Global Phishing Survey* for the second half of 2014, we found that phishing took place on 146 .TH (Thailand) domains. All 146 were on compromised servers, and 64 of those domains were on government (GO.TH) and university (AC.TH) servers.

Virtually all established TLDs experience compromised domains over time, simply because domains in those TLDs support web sites and some of them are eventually broken into. So, the chance that a TLD will experience some phishing in it is a matter of several variables, among them being the number of hosted sites in the TLD, the security of the hosters being used to support those domains, and time. Given enough time and sites, it is perhaps inevitable that phishing will occur in a TLD given the generally vulnerable state of web hosting worldwide. A corollary is that even zones that have registration requirements, and are not generally available to the public, experience some phishing. Examples include .COOP, .EDU, GO.TH, and .TRAVEL. The Appendix of this report contains statistics for all TLDs world-wide, including TLD size, number of attacks, and number of malicious registrations.

***Conclusion*: Because phishers mainly use compromised domains, we should expect that many new gTLDs will eventually experience at least some phishing over time. This is a function of vulnerable web hosting, and is possible in almost any TLD.**

Malicious Registrations

"Malicious registrations" are domains that we believe were registered specifically by phishers to perpetrate phishing. We flagged a domain as maliciously registered if it was reported for phishing within a very short time of being registered, and/or contained a brand name or misleading string in the domain itself, and/or was registered in a batch or in a pattern that indicated common ownership and intent. Malicious registrations can be suspended by registries or registrars with no risk of collateral damage.

[10] APWG quarterly *Phishing Activity Trends Reports*, at http://docs.apwg.org/reports/apwg_trends_report_q4_2014.pdf and http://docs.apwg.org/reports/apwg_trends_report_q3_2014.pdf

[11] See the APWG quarterly *Phishing Activity Trends Reports*

In 2014, we identified 49,698 malicious registrations. This is a very small percentage of the approximately 84.6 million domains created worldwide during 2014 – about 1 in 1,702 domain creates.

Of those 49,698 malicious registrations, 41,959 (84%) were registered to phish Chinese targets—services and sites in China that serve a primarily Chinese customer base. We assume that the phishers responsible are primarily Chinese.[12] Their major targets were Taobao.com, the Industrial and Commercial Bank of China (ICBC), the Bank of China (BOC), and Alipay. Since we began tracking this phenomenon in 2011, we have seen that Chinese phishers have preferred to register domains, relying upon hacked domains and compromised Web servers far less often than phishers elsewhere.

The domain strings registered by Chinese phishers are usually composed of random characters, and have no meaning in any language, such as:

hsyetf.cc, hsypar.cc, hsypaw.cc,
hsyknd.com, hsyknq.com, hvcwe.com, hyknr.com, hykx28.com

In 2014, phishers made malicious registrations in 123 TLDs total. The TLDs of choice have shifted over time, depending especially on price and sales specials. The TLDs with the most malicious registrations in 2H2014 were:

Malicious Registrations:
2H2014 leaders versus previous periods

TLD	2H2014	1H2014	2H2013	1H2013	2H2012	1H2012
com	17,018	13,623	12,347	6,477	3,145	2,598
tk	3,335	2,533	5,016	2,801	1,101	3,939
pw	1,676	2,312	860	94	0	0
cn	834	90	519	165	16	11
net	667	815	740	560	247	209
cf	626	1,282	558	0	0	0
info	474	212	763	655	516	232
ga	285	270	479	0	0	0
xyz	271	0	0	0	0	0
cc	248	26	126	16	5	3
ml	245	520	392	0	0	0

[12] These phishing attacks were advertised via e-mail lures written in Chinese, via SMS messages in Chinese sent to mobile phone customers in China, and via instant message clients popular in China such as Tencent QQ. Many of the domain registrations bear contact details with addresses in China or using Chinese service providers, and were made at Chinese registrars. Other factors about these attacks also point to perpetrators in China.

TLD	2H2014	1H2014	2H2013	1H2013	2H2012	1H2012
eu	236	99	29	53	23	7
uk	206	206	100	61	35	28
org	164	236	147	226	111	80
biz	95	28	85	52	18	5
in	66	64	46	75	180	474
co	62	41	23	26	12	11
ru	59	13	17	78	20	8
fr	44	49	146	31	9	2

In any given year, 80% to 90% of all malicious phishing registrations occur in just five to eight TLDs. As can be seen from the above chart and the Appendix file, volumes of malicious registrations have shifted from TLD to TLD over time, in response to low prices or rebates, lack of response by the registry or a particular registrar, and other factors. Examples of waves of phishing due to low prices include .CN in 2007[13], and in 2014 when the .CF, .GA, and .ML registries attracted phishers because domains in these ccTLDs were offered for free by their new operator.[14]

Phishing can also become more concentrated in a TLD due to other factors. In 2H2014 .COM contained 41.3% of the domains in the world, but 58% of the domains used for phishing. The latter percentage has grown over the last several years in part because it has been pushed up by registrations made by Chinese phishers.[15]

It is intuitive that open TLDs should be more susceptible to malicious registrations, while TLDs with registration restrictions should be less susceptible to malicious registrations. Of the 20 TLDs listed above, 18 are open, without restrictions. (.CN has nexus and documentation restrictions, and .EU and .FR are open to entities anywhere in the European Union.)

Subdomain Registries

Phishers also register subdomains at subdomain registries. These services allow users to create registrations at the third level beneath a second-level domain name that the service provider owns. [16] These services offer users their own DNS space—and often offer free DNS management – that functions

[13] http://www.apwg.com/reports/APWG_GlobalPhishingSurvey2007.pdf pages 13-14
[14] http://docs.apwg.org/reports/APWG_Global_Phishing_Report_1H_2014.pdf
[15] http://docs.apwg.org/reports/APWG_Global_Phishing_Report_2H_2014.pdf pp. 15-16; http://docs.apwg.org/reports/APWG_GlobalPhishingSurvey_2H2013.pdf pp 11-12, 14-15
[16] Examples of such services include afraid.org, the subdomain registries operated by CentralNIC, and freeavailabledomains.com which offers third-level registrations under usa.cc

just like a domain registered in a TLD registry. Thus a customer will obtain a hostname to use for his/her own Web site and/or e-mail of the form:

<customer_chooses_string>.<service_provider_sld>.TLD

For example: http://fatcebook.0zed.com

We know of more than 800 such services. One prominent subdomain registry, POPNIC.COM, claims more than 8 million registrations. (Which is more registrations than in any ccTLD registry except for .DE, .UK, and TK.)

The use of subdomain registries by phishers continues to be a challenge, because many of these registries provide the services for free, offer anonymous registration without WHOIS service, and only the subdomain providers themselves can effectively mitigate these phish.[17] While many of these services are responsive to complaints, their proactive measures to keep criminals from abusing their services are usually limited. The list of subdomain service providers used by phishers ebbs and flows over time, and each year we see phishers seeking new providers that they can exploit.

	1H2012	2H2012	1H2013	2H2013	1H2014
Second-level domains used for phishing at subdomain providers	914	773	270	795	678
Subdomains used for phishing, registered at subdomain providers	13,109	7,798	6,465	17,703	16,479
Phishing *attacks* using subdomains, registered at subdomain providers	13,307	8,294	7,134	17,678	16,986

IP Addresses

Phishing does not require the use of domain names at all. The destination can be an URL on an IP address, which a browser will resolve. An example is:

[17] Standard domain name registrars or registry operators usually cannot mitigate these phish by suspending the main or "parent" domains, because doing so would also suspend every subdomain beneath the parent, thereby affecting innocent users as well. If extensive abuse happens on a single domain, a registrar may still opt to suspend the domain based on numerous complaints. This has been observed on occasion.

http://101.55.120.53/images/login.html/

Over the past three years, 2.4% of all phishing attacks conducted worldwide have been on IP addresses rather than domains:

	1H2012	2H2012	1H2013	2H2013	1H2014	2H2014	Totals 2012-2014
Total phishing attacks	93,462	123,476	72,758	115,565	123,741	123,973	**652,975**
IP-based phishing attacks	2,419	2,489	1,972	2,394	2,891	3,582	**15,747**
Unique IP addresses used	1,864	1,841	1,626	837	2,317	3,095	**11,580**

Domain Choice and URL Construction by Phishers

By learning how phishers construct URLs, we can learn how they fool Internet users. From a technical point of view, all domain names basically work the same. A domain name in one TLD can be used for similar technical purposes as a domain in another TLD.

When using compromised domains, phishers don't have a choice of domain string – they simply use whatever domains are on the servers they break into. On compromised domains, the phisher may place a deceptive string in a subdomain or subdirectory. This can fool a user into thinking that he or she is on a trusted domain. Below are some real phishing URLs observed in late 2014.

This phishing URL on the compromised domain RACK.NO contains a misleading string in a subdirectory:

http://www.rack.no/~milit/**apple**.login1/

The below phishing URL on the compromised domain DKFMED.CA contains a misleading string (including the TLD string) in subdomains at the seventh and eighth levels:

http://**paypal.com**.security.cig-bni.scr-cdm.infocom.dkfmed.ca/

APWG reports show that the URL contains some form of the target name only between 43% and 64% of the time, depending on the period.[18] *So an URL that conveys no meaning at all, and makes no attempt to fool users, is used about one-third to one-half of the time.*

We examine all the maliciously registered domains to see if they contained a relevant brand name or reasonable variations thereof, including liberal misspellings. [19] The theory was that phishers would prefer to register domain strings that have some meaning and will help fool Internet users. Instead, we have consistently found that phishers usually do not choose (second-level) domain strings relevant to their targets. In 2012-2014, less than 9% of malicious domain registrations contained a brand name or reasonable variation thereof.

[18] APWG quarterly *Phishing Activity Trends Reports*, at
http://docs.apwg.org/reports/apwg_trends_report_q3_2014.pdf and
http://docs.apwg.org/reports/apwg_trends_report_q2_2014.pdf

[19] Examples of domain names we counted as containing brand names and therefore deceptive included: paypcil.co (PayPal), faceboook.com (Facebook), and taobaotuikkhh1.com (taobao.com)

	1H2012	2H2012	1H2013	2H2013	1H2014	2H2014	Totals 2012-2014
Domains used for phishing	64,204	89,748	53,685	82,163	87,901	95,322	**473,023**
Maliciously registered domains	7,712	5,833	12,173	22,831	22,679	27,253	**98,481**
Maliciously registered domains containing a brand name or misspelling	1,350	1,242	1,244	1,541	1,498	1,846	**8,721**

Some maliciously registered domains contain misleading strings such as “account” or “login” but these occur less frequently. Far more often than not, the domain registrations made by phishers consist of nonsense strings that have no semantic meaning whatsoever. This is evidence that deception via semantic meaning in the domain name is not necessary, and that many Internet users are not knowledgeable enough to be able to pick out the “base” or true domain name being used in a URL. This problem of user inexperience has been studied by browser manufacturers, who have created ways to help users identify the true domain name within the longer URL.[20]

There is an almost infinite supply of nonsense strings available in the legacy gTLDs and ccTLDs. Scammers can also create an almost infinite number of variations of a brand name, and it is difficult to defensively register all the resulting domain names.

This real URL on the malicious registration TA3ES.INFO spoofed the valid, trusted domain PAYPAL.US in a subdirectory:

http://ta3es.info/**paypal.us**.cgi.bin.webscr.login.webapps.mpp.home.billing/webapps/
231f6c3ad40f2dab8b9ffd80ad6356e9/

[20] Various browser features highlight the precise domain a browser is visiting in order to thwart attacks that rely on long, confusing addresses that can conceal the true domain that is being visited. For example, Internet Explorer automatically highlights what it considers to be the domain of the site the user is currently viewing. “This helps users identify the real site they’re on when a website attempts to deceive them.... Domain Highlighting effectively calls out what Internet Explorer 8 recognizes as the owning domain for the purposes of making security decisions.” (http://blogs.msdn.com/b/ie/archive/2008/03/11/address-bar-improvements-in-internet-explorer-8-beta-1.aspx). At one point Google considered not displaying URLs at all in the Chrome browser, which was a controversial idea because it would not allow savvy Internet users to recognize phishing and other deceptions: http://arstechnica.com/security/2014/05/address-bar-tweak-in-early-version-of-chrome-puts-even-savvy-users-at-risk/

The following real phishing URL , on the maliciously registered domain ABC1234.INFO, contained the misleading string “itunes.apple”:

http://abc1234.info/contact/apple.itunes/**itunes.apple**/f73c93acae9bba9ae12dff3efc854322/

.APPLE is a new gTLD. The above phish occurred in October 2014, and the .APPLE TLD has not been delegated as of this writing in 2015. Still, it is an example of how a trusted name and its TLD string can appear in a phishing URL.

***Conclusion:* A phisher does not need to register a domain in a sensitive string TLD in order to get the semantic benefit. He can simply create a URL string that *appears* to be in the sensitive TLD. Since these misleading strings appear in subdomains and subdirectories, a phisher can use any domain in any TLD to spoof any TLD string, including a sensitive TLD string.**

***Conclusion:* Phishers and other scammers can successfully leverage the meaning associated with a TLD string only if the potential victim:**

1. **can see the URL in the phishing email or in their web browser window, *and***
2. **can then correctly recognize the domain name and its TLD string within that URL, *and***
3. **places trust in the TLD string in question.**

Internationalized Domain Names

Over the past eight years, IDNs have been available at the second and third levels in many registries. IDN TLDs have gained popularity over the past three years, and allow the entire domain name to be in non-Latin characters, including the TLD extension. From 2012 through 2014,a total of 580 IDN domains were used for phishing, virtually all of them compromised domains.

Data shows that the unique characteristics of Internationalized Domain Names (IDNs) are not being used to facilitate phishing in any meaningful fashion. IDNs offer phishers two ways to fool Internet users by conveying semantic meaning, but such use has been rare.

The first technique that can use IDNs for phishing is the *homographic attack*. In this attack, a phisher seeks to deceive Internet users by exploiting the fact that characters in different language scripts may be nearly (or wholly) indistinguishable, thereby allowing the phisher to spoof a brand name.   For example, this one uses an accented letter “i”:

xn--nterbank-b2a.com → ínterbank.com

From January 2007 to December 2014 we found only nine true homographic phishing attacks, out of the hundreds of millions of domain names registered in that span.

The second technique is to convey deceptive meaning in the native language. In the second half of 2014 we found seven Chinese IDNs where the domain strings themselves were misleading, but did not attempt to exactly copy domain names owned by the targets; for example:

xn--czr93rq40bruk5heszb.com   → 工商银行首页.com   = “ICBC Home”

Although IDNs have been widely available for years, we believe that phishers have not utilized them more often because:

1. Phishers evidently don’t need to resort to such attacks – they have plenty of other, simpler options. This is also additional evidence that the domain name itself usually does not matter to a phisher.
2. By default, some browser manufacturers show the punycode version of the domain name (such as "xn--hotmal-t9a.net") in the address bar, instead of the native-character version. Users of those browsers usually can’t see homographic attacks.

Phishing Success Rates

Using the above methods, how successful are phishers, and what does this tell us about user behavior?

Browser Warnings

Many potential phishing victims are warned that they are about to visit a phishing site. Some of those users fall victim after bypassing these warnings.

Web browsers display warnings to users when a phishing attack might be occurring, based on phishing URLs provided by security companies and sophisticated detection methods. The browser will show a warning page that discourages the user from continuing. A user must click through the warning to dismiss it and proceed with her or his original task. These interstitial warning pages are not displayed for all phishing URLs, only the ones that the browser manufacturer knows about.

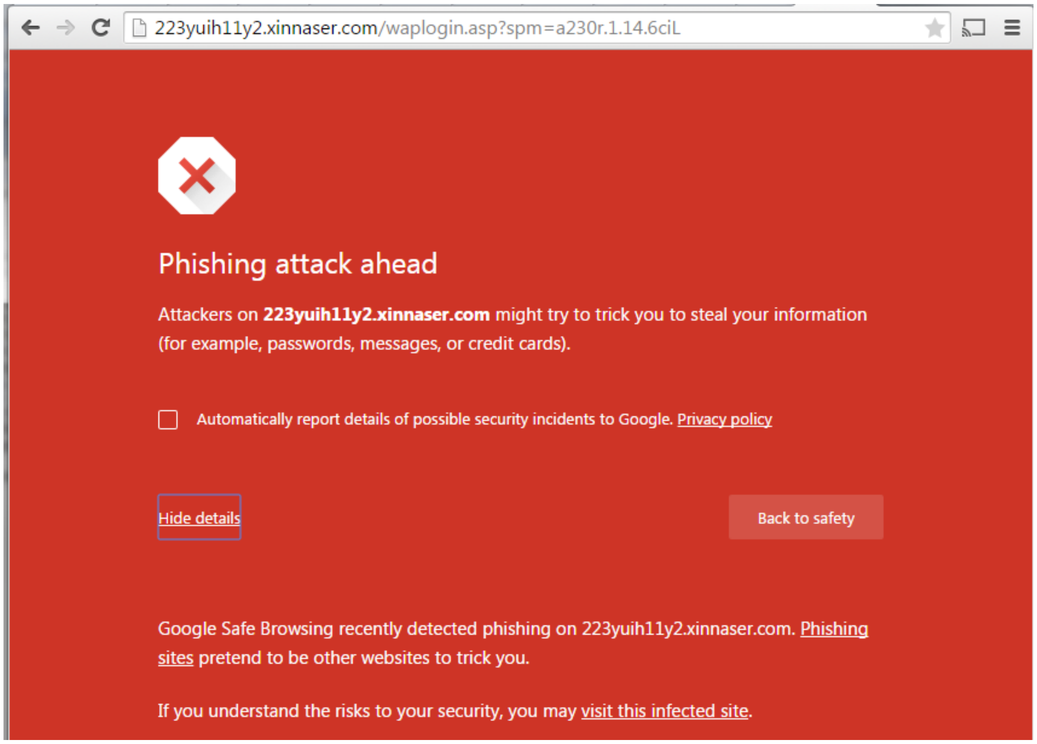

Phishing attack ahead

Attackers on **223yuih11y2.xinnaser.com** might try to trick you to steal your information (for example, passwords, messages, or credit cards).

□ Automatically report details of possible security incidents to Google. Privacy policy

Google Safe Browsing recently detected phishing on 223yuih11y2.xinnaser.com. Phishing sites pretend to be other websites to trick you.

If you understand the risks to your security, you may visit this infected site.

Google Chrome phishing warning page

The most relevant study about current browser warning pages was published in August 2013, and was entitled “Alice in Warningland: A Large-Scale Field Study of Btrowser Security Warning Effectiveness.”[21] The authors studied what users do after they are shown security warnings by their Web browsers, drawing from data collected by Google. The study encompassed 486,354 phishing attack warning impressions displayed in Google Chrome and Mozilla Firefox in May and June 2013. Users clicked through 9.1% of the time in Mozilla Firefox and 18.0% of the time in Google Chrome. So, a significant percentage of users bypassed explicit warnings.

Click Rates

Studying spear-phishing in 2014, Verizon’s *2015 Data Breach Report*[22] found that *23% of recipients opened phishing messages and 11% clicked on (malware) attachments*. This was slightly higher than in previous years, in which the rates were between 10 and 20%. So clearly the number of users who succumb to phishing attacks is great enough to support phishers.

Identity Theft Success Rates

A thorough recent study of phishing success rates is “Handcrafted Fraud and Extortion: Manual Account Hijacking in the Wild”[23] published by researchers from Google and the University of California, San Diego, in November 2014. This study found that *13.7% of visitors to phishing web forms shared their personal data with the phishers.* Some phishing sites were more successful than others -- some phishing pages were filled in by only 3% of the victims who arrived at them, while one phishing site had a 45% success rate.

***Conclusion:* To become victims, Internet users must click on an email lure, then they often click through a browser warning page, and then they must submit their data on the phishing site. In the end, human fallibility is what makes phishing possible.**

[21] Devdatta Akhawe and Adrienne Porter Felt: “Alice in Warningland: A Large-Scale Field Study of Browser Security Warning Effectiveness,” presented at 22nd USENIX Security Symposium, August 2013. https://www.usenix.org/conference/usenixsecurity13/technical-sessions/presentation/akhawe

[22] http://news.verizonenterprise.com/2015/04/2015-data-breach-report-info/

[23] Bursztein, Savage et al, “Handcrafted Fraud and Extortion: Manual Account Hijacking in the Wild,” presented at the ACM Internet Measurement Conference, November 2104 in Vancouver, British Columbia. http://conferences2.sigcomm.org/imc/2014/papers/p347.pdf

Phishing Uptimes

The "uptimes" or "live" times of phishing attacks are a vital measure of how damaging phishing attacks are, and are a metric of the success of mitigation efforts. The first day of a phishing attack is the most lucrative for the phisher as potential victims open the lure a-mails and attention is directed to the phishing attack site. So, quick takedowns of phishing sites are essential. Long-lived phish can skew the averages since some phishing sites last weeks or even months, so medians are also a useful barometer of overall mitigation efforts.

Over the past several years, the average uptimes of phishing attacks have been as follows:[24]

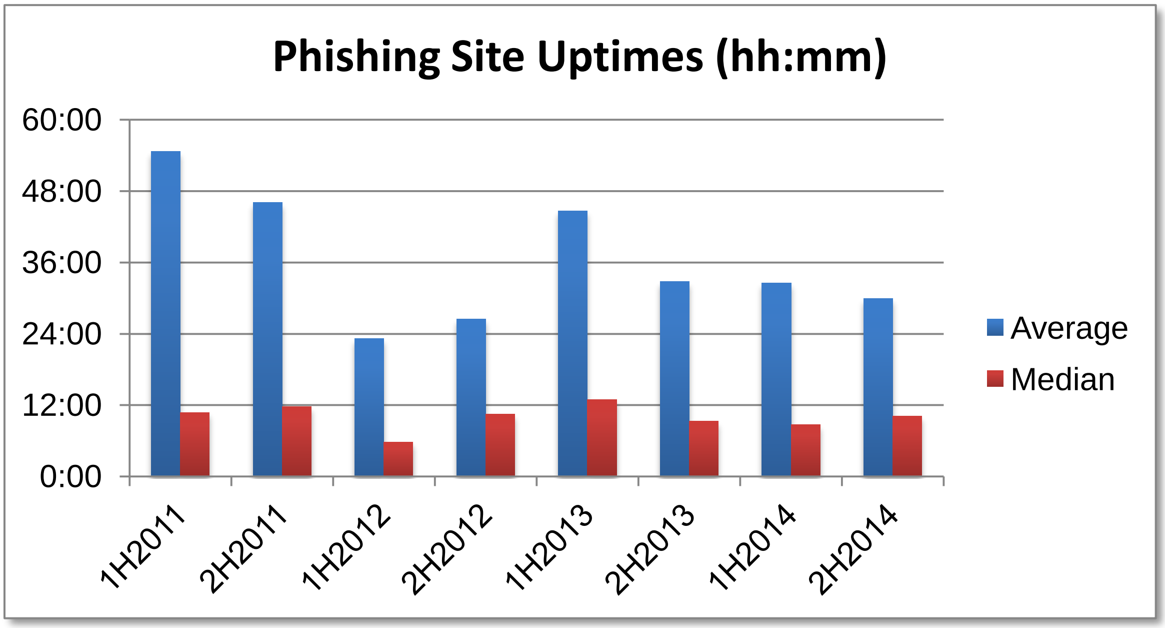

Detection and mitigation efforts by hosting providers, registries, and registrars determine how long phishing sites stay up, and therefore how much harm phishing attacks cause.

[24] Internet Identity's system tracks the uptimes automatically by monitored the phishing pages themselves. Monitoring begins as the system became aware of each phish via feeds or honeypots. Each phish was checked several times per hour to confirm its availability, and was not declared "down" until it had stayed down for at least one hour. 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 incidents and should be fair for relative comparisons.

New gTLD Analysis

As of December 2014, the new gTLDs had less phishing relative to the legacy gTLDs and ccTLDs. However, we predict that:

1) As time passes the new gTLDs may have more phishing in them than they did in 2014. Over time the new gTLDs will have more active Web sites in them, and as discussed above, some of those domains will inevitably suffer compromised hosting.
2) Some new gTLDs will attract malicious phishing registrations, just as certain legacy gTLDs and ccTLDs have attracted phishing in the past.
3) Due to these phenomena, the new gTLDs may eventually have as much phishing in them *proportionately* as in the legacy gTLDs and ccTLDs.

From 1 July to 31 December 2014:

- About 295 new gTLDs opened for registration by the public. As of 31 December, 3,684,316 domains had been registered in all new gTLDs.[25]
- Phishing occurred in 56 of those new gTLDs; 239 had no phishing at all.
- Of those 56 new gTLDs, 23 had malicious registrations made in them, often just one or two. Thirty-three had compromised domains used for phishing, often just one or two.
- A total of 454 new gTLD domain names were used for phishing. Of those, 330 were maliciously registered.
- Almost two-thirds of the phishing in the new gTLDs—288 domains—were concentrated in the .XYZ registry, which was also the largest of the new gTLDs. Of the 330 maliciously registered domains in the new gTLDs, 271 were in .XYZ.

The *APWG Global Phishing Survey* papers use the metric "Phishing Domains per 10,000", which 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. This metric is a way of revealing whether a TLD has a higher or lower incidence of phishing relative to others. In the second half of 2014, the median phishing-domains-per-10,000 score for all TLDs in the world was 3.4[26], while in the first half of 2014 it was 4.7[27].

In 2014, only nine of the 295 available gTLDs had "Phishing Domains per 10,000" scores above 3.4. It should be noted that during 2H2014, most of the new gTLDs has less than 30,000 domains in them, the minimum threshold at which the APWG reports begin to rank TLDs.

When putting compromised domains aside and considering malicious phishing registrations only, .XYZ had a significantly higher incidence of *malicious* domain registrations per 10,000 than other TLDs of all types—coming in with a score of 3.4 compared to the benchmark .COM score of 1.4, and an average of 0.9 for all TLDs[28].

[25] As per nTLDstats.com
[26] http://docs.apwg.org/reports/APWG_Global_Phishing_Report_2H_2014.pdf, page 12
[27] http://docs.apwg.org/reports/APWG_GlobalPhishingSurvey_1H2014.pdf, pages 10-11
[28] http://docs.apwg.org/reports/APWG_GlobalPhishingSurvey_2H2014.pdf, pages 11-15 and appendix

The TLD market is now more crowded and competitive than at any time in history, and some new registries are competing aggressively on price. We expect to see some gTLDs drop their prices lower than the prices for legacy gTLDs and ccTLDs, and that may attract phishing and other kinds of abuse to those new gTLDs. We note that this report examines only certain types of abuse. New gTLD domains are also being registered and used for other types of abuse, and the associated numbers and risks are beyond the score of this study.

Phishing in New gTLDs, 2014

TLD	# Unique Phishing attacks 2H2014	Unique Domain Names used for phishing 2H2014	Domains in registry, Dec 2014	# Total Malicious Domains Registered 2H2014	# Unique Phishing attacks 1H2014	Unique Domain Names used for phishing 1H2014	Domains in registry, April 2014	# Total Malicious Domains Registered 1H2014
academy	1	1	15,169					
agency	1	1	16,459		1	1	3,981	1
bayern	0	0	25,555					
berlin	3	3	155,122	1				
best	1	1	1,052					
bid	1	1	2,718	1				
bike	1	1	13,900	1				
cab	1	1	3,591					
center	4	4	27,619	3	1	1	13,939	1
cheap	4	4	3,992					
click	0	0	10,413					
club	25	22	160,591	6	3	3	1,819	1
codes	2	1	3,840					
company	4	3	35,948	2	1	1	16,614	
cruises	1	1	2,038					
dance	1	1	3,475					
diamonds	1	1	4,042					
directory	2	2	21,072					
domains	2	2	7,281					
education	2	2	13,726					
email	5	5	46,310	3	3	3	25,979	1
expert	0	0	25,843					
farm	1	1	5,878					
gallery	0	0	15,880		1	1	10,404	
guru	15	15	78,959	14	2	2	53,195	
help	2	1	2,995	1				

TLD	# Unique Phishing attacks 2H2014	Unique Domain Names used for phishing 2H2014	Domains in registry, Dec 2014	# Total Malicious Domains Registered 2H2014	# Unique Phishing attacks 1H2014	Unique Domain Names used for phishing 1H2014	Domains in registry, April 2014	# Total Malicious Domains Registered 1H2014
host	9	1	2,473					
institute	3	2	6,511					
international	3	3	15,372	1				
land	1	1	15,001		2	2	10,831	
limo	2	1	3,180					
link	18	15	53,102	7				
london	0	0	55,149					
management	1	1	8,604	1				
marketing	4	4	11,209					
media	2	2	11,602					
menu	3	2	7,201	1				
ninja	5	5	24,311					
nyc	0	0	65,361					
onl	1	1	3,719	1				
ovh	0	0	56,056					
partners	1	1	2,964	1				
photography	3	2	50,393					
photos	5	5	17,136		1	1	10,274	
pink	1	1	11,960					
pub	1	1	4,623					
qpon	1	1	482					
realtor	0	0	94,261					
report	1	1	2,907					
rocks	0	0	30,058					
ruhr	1	1	4,125	1				
sexy	3	3	17,645					
solutions	5	5	32,058	1				
support	7	6	13,383	6				
systems	1	1	14,425					
tips	2	2	33,873	1	1	1	20,991	1
today	6	6	44,025		1	1	21,890	
tokyo	0	0	30,584					
tools	1	1	5,825					
top	0	0	37,502					
training	2	2	13,372					
wang	12	10	97,591	3				
website	2	2	37,113	2				

TLD	# Unique Phishing attacks 2H2014	Unique Domain Names used for phishing 2H2014	Domains in registry, Dec 2014	# Total Malicious Domains Registered 2H2014	# Unique Phishing attacks 1H2014	Unique Domain Names used for phishing 1H2014	Domains in registry, April 2014	# Total Malicious Domains Registered 1H2014
wiki	1	1	11,130					
wtf	1	1	3,441	1				
xn--3ds443g	0	0	36,632					
xn--55qx5d	0	0	45,634					
xn--io0a7i	0	0	31,415					
xn--ses554g	0	0	107,027					
xyz	325	288	796,391	271				
zone	1	1	12,062					

Prevention and Mitigation

It is worthwhile to note that phishing *prevention* is different from phishing *mitigation*. Prevention involves lowering the chance that phishing can or will happen in the first place. Mitigation involves shutting down a phishing attempt once it is underway and has been detected. Timely mitigation can prevent users from being victimized and reduces the harm caused by phishing. As discussed above, prevention is very difficult in the case of compromised domains, because the root problem is vulnerable Web hosting, and is therefore not something that the registry operator (and the registrar, assuming it is not the hosting provider) has control over.

Malicious registrations can be reduced by controlling access to domain registrations via registration requirements, and by higher pricing. Some TLD registries and registrars have had success reducing the amount of malicious registrations in their spaces by actively suspending malicious registrations, which has acted as a deterrent.[29]

Every gTLD operator and gTLD registrar is generally able to have a terms of service that allows it to mitigate malicious registrations by suspending domain names registered for phishing and other malicious purposes. Some registries have reduced the uptimes of phishing sites by reporting compromised domains to their registrars. Detection, deterrence, quick mitigation, and anti-fraud measures can help reduce both the number of malicious registrations, and also the uptimes of all kinds of phishing attacks. In the end, security involves the balancing of risk with the various costs involved.

[29] An example is .TK and the other ccTLD operated by Freenom; see http://docs.apwg.org/reports/APWG_GlobalPhishingSurvey_2H2014.pdf paged 16-17, and .TK's malicious registration numbers in 2012-2014. The industry's largest registrar, GoDaddy, sponsors proportionately fewer gTLD domains used for phishing than some other registrars (see http://docs.apwg.org/reports/APWG_GlobalPhishingSurvey_2H2014.pdf, page 18).

Registries and registrars make choices about domain availability, price and profit, abuse, and other factors in relation to each other.[30]

> ***Conclusion:*** **To mitigate phishing, domain registries and registrars must be able to detect phishing attacks, and must be willing to suspend malicious registrations.**

[30] For a study of the incentives of "Internet intermediaries" such as domain registrars and ISPs, see Van Eeten, Michel J.G. and Johannes M. Bauer: "Economics of Malware: Security Decisions, Incentives and Externalities." Technical Report OECD STI Working Paper 2008/1: 26-34 and 46-51. Organisation for Economic Co-operation and Development, Paris.

Related Scams

Criminals perpetrate other kinds of scams that require domain names for web sites and use “social engineering” techniques similar to those used by phishers. Artists Against 419 (www.aa419.org) operates a large repository of fraud sites. Many are “419” or advance fee fraud[31] sites, while others are web sites for convincing-looking but completely fictitious banks and brokerages, non-existent shipping and escrow companies, and so on, all designed to extract money from unwary visitors in various ways.

These scammers register domains and advertise them in email lures and on social media. The domain names consist of fictitious company names, such as Transonlines Bank at Transonlines.com, and The Bureau of Diplomatic Courier & Security (a “freight forwarding department of the US department of States” [sic]) at bdcs-usa.org.

We examined the 300 most recently listed domains at Artists Against 419. Eighty percent of them used .COM domains, with 15% using .NET, and the rest using .ORG, .UK, .EU, and .INFO domains. These types of fraudsters prefer to use .COM, probably because .COM is ubiquitous and is used internationally by many legitimate businesses. The second-level domains attempt to convey some meaning, but since they are the names of bogus enterprises they are not familiar to visitors. Some uneducated visitors may trust these legitimate-sounding companies out of hand, while others are induced to trust them through web site text, promises in email, or because of persuasive phone conversations with the scammers.

Conclusion: **These scammers may use new gTLD domains in the future, but will have difficulty registering in new gTLDs that have registration requirements. Scammers of this type currently use legacy gTLDs, and those domain spaces seem to offer an ample supply of domains names for the criminals’ purposes.**

[31] For background see http://www.interpol.int/Crime-areas/Financial-crime/Fraud/419-fraud and http://en.wikipedia.org/wiki/419_scams

Acknowledgments

Thanks to Rod Rasmussen of Internet Identity, the co-author of the *Global Phishing Survey* series; and Peter Cassidy of the Anti-Phishing Working Group for his continued support of our research activities.

About the Author

Greg Aaron is President of Illumintel Inc., which provides operations and policy advising and security services to Internet companies and domain registry operators.   Mr. Aaron is an authority on the misuse of domain names and DNS policy, and regularly makes presentations at security conferences around the world. He serves as the Anti-Phishing Working Group's Senior Research Fellow, and is Co-Chair of the APWG's Internet Policy Committee. He is the co-author of the APWG's ongoing *Global Phishing Survey* series, which is the authoritative source of worldwide phishing metrics and analysis. He is a member of ICANN's Security and Stability Advisory Committee (SSAC), which provides advice and risk analysis relating to the security and integrity of the Internet's naming and address allocation systems. He was the Chair of the ICANN GNSO's Registration Abuse Policy Working Group (RAPWG). Mr. Aaron was the senior industry expert on the Ernst & Young team that evaluated new gTLD applications for ICANN in 2012-2013, reviewing the applicant responses regarding registry services, rights protection mechanisms, security, and registry operations.   Previously he worked at companies including Afilias, Travelocity, and CitySearch. Mr. Aaron is a magna cum laude graduate of the University of Pennsylvania.

Appendix: Phishing Statistics by TLD, 2012-2014

Appendix data is attached as an Excel file: "Appendix - phishing 2012-2014.xlsx"

To measure the prevalence of phishing in a TLD, APWG uses the metrics "Phishing Domains per 10,000" and "Phishing Attacks per 10,000." "Phishing Domains per 10,000"[32] 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. This metric is a way of revealing whether a TLD has a higher or lower incidence of phishing relative to others.

The metric "Phishing Attacks per 10,000" is another useful measure of the pervasiveness of phishing in a namespace. It especially highlights what TLDs are predominantly used by phishers who use subdomain services, and where high-volume phishers place multiple phish on one domain.

In 2H2014:

- The median phishing-domains-per-10,000 score was 3.4 (versus 4.7 in 1H2014).
- .COM, the world's largest and most ubiquitous TLD, had a domains-per-10,000 score of 4.7. .COM contained 58% of the phishing domains in our data set, and 41.3% of the domains in the world.
- We therefore suggest that domains-per-10,000 scores between 3.4 and 4.7 occupy the middle ground, with scores above 4.7 indicating TLDs with increasingly prevalent phishing.

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